

SANDER VAN VUGT

Cert Guide

Learn, prepare, and practice for exam success



Red Hat RHCSA 8 EX200

PEARSON IT
CERTIFICATION

<https://t.me/learningnets>

Red Hat RHCSA™ 8 Cert Guide

EX200

Sander van Vugt



Red Hat RHCSA™ 8 Cert Guide: EX200

Copyright © 2020 by Pearson Education, Inc.

All rights reserved. No part of this book shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

ISBN-13: 978-0-13-593813-3

ISBN-10: 0-13-593813-9

Library of Congress Control Number: 2019948460

ScoutAutomatedPrintCode

Trademarks

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Pearson IT Certification cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Red Hat and RHCSA are registered trademarks of Red Hat, Inc. in the United States and other countries.

Pearson IT Certification and Sander Van Vugt have no affiliation with Red Hat, Inc. The RED HAT and RHCSA trademarks are used for identification purposes only and are not intended to indicate affiliation with or approval by Red Hat, Inc.

Warning and Disclaimer

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an “as is” basis. The author and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book.

Special Sales

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

Editor-in-Chief

Mark Taub

Acquisitions Editor

Denise Lincoln

Development Editor

Ellie Bru

Managing Editor

Sandra Schroeder

Senior Project Editor

Tonya Simpson

Copy Editor

Bill McManus

Indexer

Cheryl Lenser

Proofreader

Gill Editorial Services

Technical Editor

William “Bo” Rothwell

Publishing Coordinator

Cindy Teeters

Cover Designer

Chuti Prasertsith

Compositor

codeMantra

Contents at a Glance

Introduction xxiv

Part I: Performing Basic System Management Tasks

CHAPTER 1	Installing Red Hat Enterprise Linux	3
CHAPTER 2	Using Essential Tools	29
CHAPTER 3	Essential File Management Tools	53
CHAPTER 4	Working with Text Files	81
CHAPTER 5	Connecting to Red Hat Enterprise Linux 8	99
CHAPTER 6	User and Group Management	121
CHAPTER 7	Permissions Management	143
CHAPTER 8	Configuring Networking	167

Part II: Operating Running Systems

CHAPTER 9	Managing Software	197
CHAPTER 10	Managing Processes	233
CHAPTER 11	Working with Systemd	253
CHAPTER 12	Scheduling Tasks	271
CHAPTER 13	Configuring Logging	285
CHAPTER 14	Managing Storage	309
CHAPTER 15	Managing Advanced Storage	339

Part III: Performing Advanced System Administration Tasks

CHAPTER 16	Basic Kernel Management	367
CHAPTER 17	Managing and Understanding the Boot Procedure	385
CHAPTER 18	Essential Troubleshooting Skills	401
CHAPTER 19	An Introduction to Bash Shell Scripting	423

Part IV: Managing Network Services

CHAPTER 20	Configuring SSH	443
CHAPTER 21	Managing Apache HTTP Services	457
CHAPTER 22	Managing SELinux	471

CHAPTER 23	Configuring a Firewall	497
CHAPTER 24	Accessing Network Storage	511
CHAPTER 25	Configuring Time Services	527
CHAPTER 26	Final Preparation	541
CHAPTER 27	Theoretical Pre-Assessment Exam	545

Part V: RHCSA RHEL 8 Practice Exams

	RHCSA Practice Exam A	549
	RHCSA Practice Exam B	551
APPENDIX A	Answers to the “Do I Know This Already?” Quizzes and Review Questions	553
	Glossary	587
	Index	609

Online Elements:

	RHCSA Practice Exam C
	RHCSA Practice Exam D
APPENDIX B	Memory Tables
APPENDIX C	Memory Tables Answer Key
APPENDIX D	Study Planner
	Glossary

Table of Contents

Introduction xxiv

Part I: Performing Basic System Management Tasks

Chapter 1 Installing Red Hat Enterprise Linux 3

“Do I Know This Already?” Quiz 3

Foundation Topics 6

Preparing to Install Red Hat Enterprise Linux 6

What Is Red Hat Enterprise Linux 8 Server? 6

Getting the Software 7

Using Red Hat Enterprise Linux 7

Using CentOS 7

Other Distributions 8

Understanding Access to Repositories 8

Setup Requirements 9

Cert Guide Environment Description 9

Performing a Manual Installation 10

Summary 24

Exam Preparation Tasks 25

Review Questions 25

End-of-Chapter Lab 26

Chapter 2 Using Essential Tools 29

“Do I Know This Already?” Quiz 29

Foundation Topics 32

Basic Shell Skills 32

Understanding Commands 32

Executing Commands 32

I/O Redirection 34

Using Pipes 36

History 37

Bash Completion 39

Editing Files with vim 39

Understanding the Shell Environment	42
Understanding Variables	42
Recognizing Environment Configuration Files	43
Using /etc/motd and /etc/issue	43
Finding Help	44
Using --help	44
Using man	45
Finding the Right man Page	46
Updating mandb	47
Using info	48
Using /usr/share/doc Documentation Files	49
Summary	50
Exam Preparation Tasks	50
Review Questions	51
End-of-Chapter Lab	51
Chapter 3 Essential File Management Tools	53
“Do I Know This Already?” Quiz	53
Foundation Topics	56
Working with the File System Hierarchy	56
Defining the File System Hierarchy	56
Understanding Mounts	57
Managing Files	61
Working with Wildcards	62
Managing and Working with Directories	62
Working with Absolute and Relative Pathnames	63
Listing Files and Directories	64
Copying Files	65
Moving Files	66
Deleting Files	66
Using Links	68
Understanding Hard Links	68
Understanding Symbolic Links	69
Creating Links	69
Removing Links	70

	Working with Archives and Compressed Files	71
	Managing Archives with tar	72
	<i>Creating Archives with tar</i>	72
	<i>Monitoring and Extracting tar Files</i>	73
	Using Compression	74
	Summary	75
	Exam Preparation Tasks	75
	Review Questions	77
	End-of-Chapter Lab	77
Chapter 4	Working with Text Files	81
	“Do I Know This Already?” Quiz	81
	Foundation Topics	84
	Using Common Text File–Related Tools	84
	Doing More with less	84
	Showing File Contents with cat	85
	Displaying the First or Last Lines of a File with head and tail	86
	Filtering Specific Columns with cut	87
	Sorting File Contents and Output with sort	87
	Counting Lines, Words, and Characters with wc	88
	A Primer to Using Regular Expressions	89
	Using Line Anchors	90
	Using Escaping in Regular Expressions	90
	Using Wildcards and Multipliers	91
	Using grep to Analyze Text	92
	Working with Other Useful Text Processing Utilities	93
	Summary	94
	Exam Preparation Tasks	95
	Review Questions	96
	End-of-Chapter Lab	96
Chapter 5	Connecting to Red Hat Enterprise Linux 8	99
	“Do I Know This Already?” Quiz	99
	Foundation Topics	102
	Working on Local Consoles	102

	Logging In to a Local Console	102
	Switching Between Terminals in a Graphical Environment	103
	Working with Multiple Terminals in a Nongraphical Environment	104
	Understanding Pseudo Terminal Devices	106
	Booting, Rebooting, and Shutting Down Systems	106
	Using SSH and Related Utilities	108
	Accessing Remote Systems Using SSH	108
	Using Graphical Applications in an SSH Environment	111
	Securely Transferring Files Between Systems	111
	<i>Using scp to Securely Copy Files</i>	112
	<i>Using sftp to Securely Transfer Files</i>	112
	<i>Using rsync to Synchronize Files</i>	113
	Configuring Key-Based Authentication for SSH	114
	Using Passphrases or Not?	114
	Using the screen Command	115
	Summary	117
	Exam Preparation Tasks	117
	Review Questions	118
	End-of-Chapter Labs	118
Chapter 6	User and Group Management	121
	“Do I Know This Already?” Quiz	121
	Foundation Topics	124
	Different User Types	124
	Users on Linux	124
	Working as Root	124
	Using su	125
	sudo	126
	PolicyKit	126
	Managing User Accounts	127
	System Accounts and Normal Accounts	127
	Creating Users	131
	<i>Modifying the Configuration Files</i>	131

	<i>Using useradd</i>	131
	<i>Home Directories</i>	132
	Managing User Properties	132
	Configuration Files for User Management Defaults	132
	Managing Password Properties	133
	Creating a User Environment	134
	Creating and Managing Group Accounts	135
	Understanding Linux Groups	135
	Creating Groups	136
	<i>Creating Groups with vigr</i>	136
	<i>Using groupadd to Create Groups</i>	137
	Managing Group Properties	137
	Summary	138
	Exam Preparation Tasks	138
	Review Questions	140
	End-of-Chapter Lab	140
Chapter 7	Permissions Management	143
	“Do I Know This Already?” Quiz	143
	Foundation Topics	146
	Managing File Ownership	146
	Displaying Ownership	146
	Changing User Ownership	147
	Changing Group Ownership	148
	Understanding Default Ownership	148
	Managing Basic Permissions	149
	Understanding Read, Write, and Execute Permissions	149
	Applying Read, Write, and Execute Permissions	150
	Managing Advanced Permissions	152
	Understanding Advanced Permissions	152
	Applying Advanced Permissions	155
	Managing ACLs	156
	Understanding ACLs	156
	Preparing Your File System for ACLs	157
	Changing and Viewing ACL Settings with setfacl and getfacl	157
	Working with Default ACLs	159

Setting Default Permissions with umask	160
Working with User-Extended Attributes	162
Summary	163
Exam Preparation Tasks	163
Review Questions	164
End-of-Chapter Lab	165
Chapter 8 Configuring Networking	167
“Do I Know This Already?” Quiz	167
Foundation Topics	170
Networking Fundamentals	170
IP Addresses	170
IPv6 Addresses	171
Network Masks	171
Binary Notation	172
MAC Addresses	173
Protocol and Ports	173
Managing Network Addresses and Interfaces	174
Validating Network Configuration	175
Validating Network Address Configuration	175
Validating Routing	178
Validating the Availability of Ports and Services	178
Configuring Network Configuration with nmtui and nmcli	180
Required Permissions to Change Network Configuration	180
Configuring the Network with nmcli	181
Configuring the Network with nmtui	187
Working on Network Configuration Files	189
Setting Up Hostname and Name Resolution	190
Hostnames	190
DNS Name Resolution	192
Summary	194
Exam Preparation Tasks	194
Review Questions	195
End-of-Chapter Lab	195

Part II: Operating Running Systems

Chapter 9 Managing Software 197

“Do I Know This Already?” Quiz	197
Foundation Topics	200
Managing Software Packages with yum	200
Understanding the Role of Repositories	200
Registering Red Hat Enterprise Linux for Support	201
Managing Subscriptions	202
Specifying Which Repository to Use	203
Understanding Repository Security	206
Creating Your Own Repository	208
Using yum	209
Using yum to Find Software Packages	210
Getting More Information About Packages	211
Installing and Removing Software Packages	212
Showing Lists of Packages	214
<i>Updating Packages</i>	215
<i>Working with yum Package Groups</i>	216
<i>Using yum History</i>	218
Managing Package Module Streams	220
Understanding Modules	220
Managing Modules	221
Managing Software Packages with rpm	225
Understanding RPM Filenames	226
Querying the RPM Database	226
Querying RPM Package Files	227
Using repoquery	228
Summary	229
Exam Preparation Tasks	229
Review Questions	230
End-of-Chapter Labs	231

Chapter 10 Managing Processes 233

- “Do I Know This Already?” Quiz 233
- Foundation Topics 236
- Introduction to Process Management 236
- Managing Shell Jobs 236
 - Running Jobs in the Foreground and Background 236
 - Managing Shell Jobs 237
 - Managing Parent-Child Relations 239
- Using Common Command-Line Tools for Process Management 239
 - Understanding Processes and Threads 239
 - Using ps to Get Process Information 240
 - Adjusting Process Priority with nice 242
 - Sending Signals to Processes with kill, killall, and pkill 244
- Using top to Manage Processes 246
- Using tuned to Optimize Performance 248
- Summary 249
- Exam Preparation Tasks 249
- Review Questions 250
- End-of-Chapter Lab 251

Chapter 11 Working with Systemd 253

- “Do I Know This Already?” Quiz 253
- Foundation Topics 256
- Understanding Systemd 256
 - Understanding Systemd Unit Locations 256
 - Understanding Systemd Service Units 257
 - Understanding Systemd Mount Units 258
 - Understanding Systemd Socket Units 259
 - Understanding Systemd Target Units 259
- Managing Units Through Systemd 261
 - Managing Dependencies 263
 - Managing Unit Options 265
- Summary 266
- Exam Preparation Tasks 267
- Review Questions 268
- End-of-Chapter Lab 268

Chapter 12 Scheduling Tasks 271

- “Do I Know This Already?” Quiz 271
- Foundation Topics 274
- Configuring Cron to Automate Recurring Tasks 274
 - Managing the crond Service 274
 - Understanding Cron Timing 276
 - Managing Cron Configuration Files 276
 - Understanding the Purpose of anacron 278
 - Managing Cron Security 280
- Configuring at to Schedule Future Tasks 281
- Summary 282
- Exam Preparation Tasks 282
- Review Questions 283
- End-of-Chapter Lab 283

Chapter 13 Configuring Logging 285

- “Do I Know This Already?” Quiz 285
- Foundation Topics 288
- Understanding System Logging 288
 - Understanding the Role of rsyslogd and journald 288
 - Reading Log Files 290
 - Understanding Log File Contents 290
 - Live Log File Monitoring 292
 - Using logger 292
- Configuring rsyslogd 293
 - Understanding rsyslogd Configuration Files 293
 - Understanding rsyslog.conf Sections 293
 - Understanding Facilities, Priorities, and Log Destinations 294
- Rotating Log Files 297
- Working with journald 299
 - Using journalctl to Find Events 299
 - Preserving the Systemd Journal 302
- Summary 304
- Exam Preparation Tasks 304
- Review Questions 305
- End-of-Chapter Lab 306

Chapter 14 Managing Storage 309

- “Do I Know This Already?” Quiz 309
- Foundation Topics 312
- Understanding MBR and GPT Partitions 312
 - Understanding the MBR Partitioning Scheme 312
 - Understanding the Need for GPT Partitioning 313
 - Understanding Storage Measurement Units 314
- Managing Partitions and File Systems 314
 - Creating MBR Partitions with fdisk 315
 - Using Extended and Logical Partitions on MBR 319
 - Creating GPT Partitions with gdisk 320
 - Creating GPT Partitions with parted 324
 - Creating File Systems 325
 - Changing File System Properties 327
 - Managing Ext4 File System Properties* 327
 - Managing XFS File System Properties* 329
 - Adding Swap Partitions 329
 - Adding Swap Files 330
- Mounting File Systems 330
 - Manually Mounting File Systems 331
 - Using Device Names, UUIDs, or Disk Labels 331
 - Automating File System Mounts Through /etc/fstab 332
- Summary 335
- Exam Preparation Tasks 335
- Review Questions 336
- End-of-Chapter Lab 337

Chapter 15 Managing Advanced Storage 339

- “Do I Know This Already?” Quiz 339
- Foundation Topics 342
- Understanding LVM 342
 - LVM Architecture 342
 - LVM Features 343
- Creating LVM Logical Volumes 344

Creating the Physical Volumes	346
Creating the Volume Groups	349
Creating the Logical Volumes and File Systems	351
Understanding LVM Device Naming	351
Resizing LVM Logical Volumes	353
Resizing Volume Groups	353
Resizing Logical Volumes and File Systems	354
Configuring Stratis	356
Understanding Stratis Architecture	356
Creating Stratis Storage	357
Managing Stratis	358
Configuring VDO	359
Understanding VDO	359
Setting Up VDO	360
Summary	362
Exam Preparation Tasks	362
Review Questions	363
End-of-Chapter Labs	364

Part III: Performing Advanced System Administration Tasks

Chapter 16 Basic Kernel Management 367

“Do I Know This Already?” Quiz	367
Foundation Topics	370
Understanding the Role of the Linux Kernel	370
Understanding the Use of Kernel Threads and Drivers	370
Analyzing What the Kernel Is Doing	371
Working with Kernel Modules	373
Understanding Hardware Initialization	374
Managing Kernel Modules	376
Checking Driver Availability for Hardware Devices	379
Managing Kernel Module Parameters	380
Upgrading the Linux Kernel	381

Summary	381
Exam Preparation Tasks	381
Review Questions	382
End-of-Chapter Lab	383
Chapter 17 Managing and Understanding the Boot Procedure	385
“Do I Know This Already?” Quiz	385
Foundation Topics	388
Managing Systemd Targets	388
Understanding Systemd Targets	388
<i>Working with Targets</i>	388
<i>Understanding Target Units</i>	389
<i>Understanding Wants</i>	390
<i>Managing Systemd Targets</i>	390
<i>Isolating Targets</i>	391
<i>Setting the Default Target</i>	393
Working with GRUB 2	394
Understanding GRUB 2	394
Understanding GRUB 2 Configuration Files	395
Modifying Default GRUB 2 Boot Options	395
Summary	396
Exam Preparation Tasks	397
Review Questions	397
End-of-Chapter Labs	398
Chapter 18 Essential Troubleshooting Skills	401
“Do I Know This Already?” Quiz	401
Foundation Topics	404
Understanding the RHEL 8 Boot Procedure	404
Passing Kernel Boot Arguments	406
Accessing the Boot Prompt	406
Starting a Troubleshooting Target	407
Using a Rescue Disk	408
Restoring System Access Using a Rescue Disk	409
Reinstalling GRUB Using a Rescue Disk	411
Re-Creating the Initramfs Using a Rescue Disk	411

Fixing Common Issues	413
Reinstalling GRUB 2	413
Fixing the Initramfs	414
Recovering from File System Issues	414
Resetting the Root Password	416
Recovering Access to a Virtual Machine	416
Summary	418
Exam Preparation Tasks	418
Review Questions	419
End-of-Chapter Lab	420
Chapter 19 An Introduction to Bash Shell Scripting	423
“Do I Know This Already?” Quiz	423
Foundation Topics	426
Understanding Shell Scripting Core Elements	426
Using Variables and Input	428
Using Positional Parameters	428
Working with Variables	430
Using Conditional Loops	432
Working with if ... then ... else	432
Using and &&	433
Applying for	434
Understanding while and until	435
Understanding case	436
Bash Shell Script Debugging	438
Summary	438
Exam Preparation Tasks	438
Review Questions	439
End-of-Chapter Lab	440
Part IV: Managing Network Services	
Chapter 20 Configuring SSH	443
“Do I Know This Already?” Quiz	443
Foundation Topics	446
Hardening the SSH Server	446

	Limiting Root Access	446
	Configuring Alternative Ports	447
	Modifying SELinux to Allow for Port Changes	447
	Limiting User Access	448
	Using Other Useful sshd Options	449
	Session Options	450
	Connection Keepalive Options	450
	Configuring Key-Based Authentication with Passphrases	451
	Summary	452
	Exam Preparation Tasks	452
	Review Questions	453
	End-of-Chapter Lab	454
Chapter 21	Managing Apache HTTP Services	457
	“Do I Know This Already?” Quiz	457
	Foundation Topics	460
	Configuring a Basic Apache Server	460
	Installing the Required Software	460
	Identifying the Main Configuration File	460
	Creating Web Server Content	463
	Understanding Apache Configuration Files	464
	Creating Apache Virtual Hosts	465
	Summary	468
	Exam Preparation Tasks	468
	Review Questions	469
	End-of-Chapter Lab	469
Chapter 22	Managing SELinux	471
	“Do I Know This Already?” Quiz	472
	Foundation Topics	475
	Understanding SELinux Working and Modes	475
	Understanding Context Settings and the Policy	479
	Monitoring Current Context Labels	479
	Setting Context Types	481
	Finding the Context Type You Need	484

	Restoring Default File Contexts	485
	Using Boolean Settings to Modify SELinux Settings	487
	Diagnosing and Addressing SELinux Policy Violations	488
	Making SELinux Analyzing Easier	490
	Preparing for SELinux Coverage on the Exam	492
	Summary	492
	Exam Preparation Tasks	492
	Review Questions	494
	End-of-Chapter Lab	494
Chapter 23	Configuring a Firewall	497
	“Do I Know This Already?” Quiz	497
	Foundation Topics	500
	Understanding Linux Firewalling	500
	Understanding Previous Solutions	500
	Understanding Firewallld	500
	Understanding Firewallld Zones	501
	Understanding Firewallld Services	502
	Working with Firewallld	503
	Summary	506
	Exam Preparation Tasks	506
	Review Questions	507
	End-of-Chapter Lab	508
Chapter 24	Accessing Network Storage	511
	“Do I Know This Already?” Quiz	511
	Foundation Topics	514
	Using NFS Services	514
	Understanding NFS Security	514
	RHEL 8 NFS Versions	514
	Mounting the NFS Share	515
	Mounting the NFS Share	516
	Using CIFS Services	516
	Discovering CIFS Shares	517
	Mounting Samba Shares	518
	Authenticating to Samba Shares	518

Configuring a Samba Server	518
Mounting Remote File Systems Through fstab	520
Mounting NFS Shares Through fstab	520
Mounting Samba Shares Through fstab	521
Using Automount to Mount Remote File Systems	521
Understanding Automount	521
Defining Mounts in Automount	522
Configuring Automount for NFS	522
Using Wildcards in Automount	523
Summary	524
Exam Preparation Tasks	524
Review Questions	525
End-of-Chapter Lab	525
Chapter 25 Configuring Time Services	527
“Do I Know This Already?” Quiz	527
Foundation Topics	530
Understanding Local Time	530
Using Network Time Protocol	531
Managing Time on Red Hat Enterprise Linux	531
Using date	532
Using hwclock	532
Using timedatectl	533
Managing Time Zone Settings	535
Configuring Time Service Clients	537
Summary	538
Exam Preparation Tasks	539
Review Questions	540
End-of-Chapter Lab	540
Chapter 26 Final Preparation	541
General Tips	541
Verifying Your Readiness	541
Registering for the Exam	541
On Exam Day	542
<i>During the Exam</i>	542
<i>The Nondisclosure Agreement</i>	544

Chapter 27 Theoretical Pre-Assessment Exam 545**Part V: RHCSA RHEL 8 Practice Exams**

RHCSA Practice Exam A 549

RHCSA Practice Exam B 551

**Appendix A Answers to the “Do I Know This Already?”
Quizzes and Review Questions 553**

Glossary 587

Index 609

Online Elements:

RHCSA Practice Exam C

RHCSA Practice Exam D

Appendix B Memory Tables**Appendix C Memory Tables Answer Key****Appendix D Study Planner**

Glossary

About the Author

Sander van Vugt is an independent Linux trainer, author, and consultant living in the Netherlands. Sander is the author of the best-selling *Red Hat Certified System Administrator (RHCSA) Complete Video Course* and the *Red Hat Certified Engineer (RHCE) Complete Video Course*. He has also written numerous books about different Linux-related topics and many articles for Linux publications around the world. Sander has been teaching Red Hat, Linux+, and LFCS classes since 1994. As a consultant, he specializes in Linux high-availability solutions and performance optimization. You can find more information about Sander on his website at <http://www.sandervanvugt.com>.

For more information about RHCSA certification and additional resources, visit the author's Red Hat Certification page at <http://www.rhatcert.com/>.

Dedication

This book is dedicated to my family: Florence, Franck, and Alex. Together we've made great accomplishments over the past year.

Acknowledgments

This book could not have been written without the help of all the people who contributed to it. I want to thank the people at Pearson, Denise Lincoln and Ellie Bru in particular. We've worked a lot together over the years, and this book is another milestone on our road to success!

About the Technical Reviewer

William “Bo” Rothwell, at the impressionable age of 14, crossed paths with a TRS-80 Micro Computer System (affectionately known as a Trash 80). Soon after, the adults responsible for Bo made the mistake of leaving him alone with the TRS-80. He immediately dismantled it and held his first computer class, showing his friends what made this “computer thing” work.

Since this experience, Bo’s passion for understanding how computers work and sharing this knowledge with others has resulted in a rewarding career in IT training. His experience includes Linux, Unix, IT security, DevOps, and programming languages such as Perl, Python, Tcl, and Bash. Bo is the founder and lead instructor of One Course Source, an IT training organization.

Bo is an author of several books, including *Linux for Developers: Jumpstart Your Linux Programming Skills*, *Linux Essentials for Cybersecurity*, and *LPIC-2 Cert Guide*. He can be reached on LinkedIn: <https://www.linkedin.com/in/bo-rothwell/>.

We Want to Hear from You!

As the reader of this book, *you* are our most important critic and commentator. We value your opinion and want to know what we’re doing right, what we could do better, what areas you’d like to see us publish in, and any other words of wisdom you’re willing to pass our way.

We welcome your comments. You can email or write to let us know what you did or didn’t like about this book—as well as what we can do to make our books better.

Please note that we cannot help you with technical problems related to the topic of this book.

When you write, please be sure to include this book’s title and author as well as your name and email address. We will carefully review your comments and share them with the author and editors who worked on the book.

Email: community@informit.com

Reader Services

Register your copy of *Red Hat RHCSA 8 Cert Guide* at www.pearsonitcertification.com for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to www.pearsonitcertification.com/register and log in or create an account*. Enter the product ISBN 9780135938133 and click Submit. When the process is complete, you will find any available bonus content under Registered Products.

*Be sure to check the box that you would like to hear from us to receive exclusive discounts on future editions of this product.

Introduction

Welcome to the *Red Hat RHCSA 8 Cert Guide*. The Red Hat exams are some of the toughest in the business, and this guide will be an essential tool in helping you prepare to take the Red Hat Certified System Administrator (RHCSA) exam.

As an instructor with more than 20 years of experience teaching Red Hat Enterprise Linux, I have taken the RHCSA exam (and the RHCE exam) numerous times so that I can keep current on the progression of the exam, what is new, and what is different. I share my knowledge with you in this comprehensive Cert Guide so that you get the guidance you need to pass the RHCSA exam.

The RHCSA exam was recently updated for Red Hat Enterprise Linux 8. This book contains all you need to know to pass the RHCSA exam. As you will see, this Cert Guide covers every objective in the exam and comprises 27 chapters, more than 80 exercises, 4 practice exams, an extensive glossary, and hours of video training. This Cert Guide is the best resource you can get to prepare for and pass the RHCSA exam.

Goals and Methods

To learn the topics described in this book, it is recommended that you create your own testing environment. You cannot become an RHCSA without practicing a lot. Within the exercises included in every chapter of the book, you will find all the examples you need to understand what is on the exam and thoroughly learn the material needed to pass it. The exercises in the chapters provide step-by-step procedures that you can follow to find working solutions so that you can get real experience before taking the exam.

Each chapter also includes one or more end-of-chapter labs. These labs ask questions that are similar to the questions that you might encounter on the exam. Solutions are not provided for these labs, and that is on purpose, because you need to train yourself to verify your work before you take the exam. On the exam, you also have to be able to verify for yourself whether the solution is working as expected. Please be sure to also go to this book's companion website, which provides additional practice exams, appendixes, and video training—all key components to studying for and passing the exam.

To make working with the assignments in this book as easy as possible, the complete lab environment is Bring Your Own. In Chapter 1 you'll learn how to install CentOS or Red Hat Enterprise Linux 8 in a virtual machine, and that is all that is required to go through the labs. You don't need to import any virtual machines; just install your own virtual machine and you'll be ready to go!

This book contains everything you need to pass the exam, but if you want more guidance and practice, I have a number of video training titles available to help you study, including the following:

- *Red Hat Certified System Administrator (RHCSA) Complete Video Course*, 3rd Edition
- *Upgrading to Red Hat Enterprise Linux (RHEL) 8 LiveLessons*

Apart from these products, you might also appreciate my website, <http://rhatcert.com>. Through this website, I provide updates on anything that is useful to exam candidates. I recommend that you register on the website so that I can send you messages about important updates that I've made available. Also, you'll find occasional video updates on my YouTube channel, rhatcert. I hope that all these resources provide you with everything you need to pass the Red Hat Certified System Administrator exam in an affordable way. Good luck!

Who Should Read This Book?

This book is written as an RHCSA exam preparation guide. That means that you should read it if you want to increase your chances of passing the RHCSA exam. A secondary use of this book is as a reference guide for Red Hat system administrators. As an administrator, you'll like the explanations and procedures that describe how to get things done on Red Hat Enterprise Linux.

So, why should you consider passing the RHCSA exam? That question is simple to answer: Linux has become a very important operating system, and qualified professionals are in demand all over the world. If you want to work as a Linux professional and prove your skills, the RHCSA certificate really helps and is one of the most sought-after certificates in IT. Having this certificate dramatically increases your chances of becoming hired as a Linux professional.

How This Book Is Organized

This book is organized as a reference guide to help you prepare for the RHCSA exam. If you're new to the topics, you can just read it cover to cover. You can also read the individual chapters that you need to fine-tune your skills in this book. Every chapter starts with a "Do I Know This Already?" quiz that asks questions about ten topics that are covered in that chapter and provides a simple tool to check whether you're already familiar with the topics covered in the chapter.

The book also provides four RHCSA practice exams; these are an essential part of readying yourself for the real exam experience. You may be able to provide the right

answer to the multiple-choice chapter questions, but that doesn't mean that you can create the configurations when you take the exam. The companion files include two extra practice exams, two hours of video from the *Red Hat Certified System Administrator (RHCSA) Complete Video Course*, 3rd Edition, and additional appendixes. The following outline describes the topics that are covered in the chapters:

Part I: Performing Basic System Management Tasks

- **Chapter 1: Installing Red Hat Enterprise Linux:** In this chapter, you learn how to install Red Hat Enterprise Linux Server (RHEL). It also shows how to set up an environment that can be used for working on the labs and exercises in this book.
- **Chapter 2: Using Essential Tools:** This chapter covers some of the Linux basics, including working with the shell and Linux commands. This chapter is particularly important if you're new to working with Linux.
- **Chapter 3: Essential File Management Tools:** In this chapter, you learn how to work with tools to manage the Linux file system. This is an important skill because everything on Linux is very file system oriented.
- **Chapter 4: Working with Text Files:** In this chapter, you learn how to work with text files. The chapter teaches you how to create text files, but also how to look for specific content in the different text files.
- **Chapter 5: Connecting to Red Hat Enterprise Linux 8:** This chapter describes the different methods that can be used to connect to RHEL 8. It explains both local login and remote login and the different terminal types used for this purpose.
- **Chapter 6: User and Group Management:** On Linux, users are entities that can be used by people or processes that need access to specific resources. This chapter explains how to create users and make user management easier by working with groups.
- **Chapter 7: Permissions Management:** In this chapter, you learn how to manage Linux permissions through the basic read, write, and execute permissions, but also through the special permissions and access control lists.
- **Chapter 8: Configuring Networking:** A server is useless if it isn't connected to a network. In this chapter, you learn the essential skills required for managing network connections.

Part II: Operating Running Systems

- **Chapter 9: Managing Software:** Red Hat offers an advanced system for managing software packages. This chapter teaches you how it works.

- **Chapter 10: Managing Processes:** As an administrator, you need to know how to work with the different tasks that can be running on Linux. This chapter shows how to do this, by sending signals to processes and by changing process priority.
- **Chapter 11: Working with Systemd:** Systemd is the standard manager of services and more in RHEL 8. In this chapter you learn how to manage services using Systemd.
- **Chapter 12: Scheduling Tasks:** In this chapter, you learn how to schedule a task for execution on a moment that fits you best.
- **Chapter 13: Configuring Logging:** As an administrator, you need to know what's happening on your server. The rsyslogd and journald services are used for this purpose. This chapter explains how to work with them.
- **Chapter 14: Managing Storage:** Storage management is an important skill to master as a Linux administrator. This chapter explains how hard disks can be organized in partitions and how these partitions can be mounted in the file system.
- **Chapter 15: Managing Advanced Storage:** Dividing disks in partitions isn't very flexible. If you need optimal flexibility, you need LVM logical volumes, which are used by default while you're installing Red Hat Enterprise Linux. This chapter shows how to create and manage those logical volumes. You'll also learn how to work with the new Stratis and VDO storage techniques.

Part III: Performing Advanced System Administration Tasks

- **Chapter 16: Basic Kernel Management:** The kernel is the part of the operating system that takes care of handling hardware. This chapter explains how that works and what an administrator can do to analyze the current configuration and manage hardware devices in case the automated procedure doesn't work well.
- **Chapter 17: Managing and Understanding the Boot Procedure:** Many things are happening when a Linux server boots. This chapter describes the boot procedure in detail and zooms in on vital aspects of the boot procedure, including the GRUB 2 boot loader and the Systemd service manager.
- **Chapter 18: Essential Troubleshooting Skills:** Sometimes a misconfiguration can cause your server to no longer boot properly. This chapter teaches you some of the techniques that can be applied when normal server startup is no longer possible.

- **Chapter 19: An Introduction to Bash Shell Scripting:** Some tasks are complex and need to be performed repeatedly. Such tasks are ideal candidates for optimization through shell scripts. In this chapter, you learn how to use conditional structures in shell scripts to automate tasks efficiently.

Part IV: Managing Network Services

- **Chapter 20: Configuring SSH:** Secure Shell (SSH) is one of the fundamental services that is enabled on RHEL 8 by default. Using SSH allows you to connect to a server remotely. In this chapter you learn how to set up an SSH server.
- **Chapter 21: Managing Apache HTTP Services:** Apache is the most commonly used service on Linux. This chapter shows how to set up Apache web services, including the configuration of Apache virtual hosts.
- **Chapter 22: Managing SELinux:** Many Linux administrators only know how to switch it off, because SELinux is hard to manage and is often why services cannot be accessed. In this chapter, you learn how SELinux works and what to do to configure it so that your services are still working and will be much better protected against possible abuse.
- **Chapter 23: Configuring a Firewall:** Apart from SELinux, RHEL 8 comes with a firewall as one of the main security measures, which is implemented by the firewalld service. In this chapter, you learn how this service is organized and what you can do to block or enable access to specific services.
- **Chapter 24: Accessing Network Storage:** While working in a server environment, managing remote mounts is an important skill. A remote mount allows a client computer to access a file system offered through a remote server. These remote mounts can be made through a persistent mount in /etc/fstab, or by using the automount service. This chapter teaches how to set up either of them and shows how to configure an FTP server.
- **Chapter 25: Configuring Time Services:** For many services, such as databases and Kerberos, it is essential to have the right time. That's why as an administrator you need to be able to manage time on Linux. This chapter teaches you how.
- **Chapter 26: Final Preparation:** In this chapter, you get some final exam preparation tasks. It contains many tips that help you maximize your chances of passing the RHCSA exam.
- **Chapter 27: Theoretical Pre-Assessment Exam:** This chapter provides an RHCSA theoretical pre-assessment exam to help you assess your skills and determine the best route forward for studying for the exam.

Part V: RHCSA RHEL 8 Practice Exams

This part supplies two RHCSA practice exams so that you can test your knowledge and skills further before taking the exams. Two additional exams are on the companion website.

How to Use This Book

To help you customize your study time using this book, the core chapters have several features that help you make the best use of your time:

- **“Do I Know This Already?” Quizzes:** Each chapter begins with a quiz that helps you determine the amount of time you need to spend studying that chapter and the specific topics that you need to focus on.
- **Foundation Topics:** These are the core sections of each chapter. They explain the protocols, concepts, and configuration for the topics in that chapter.
- **Exam Preparation Tasks:** Following the “Foundation Topics” section of each chapter, the “Exam Preparation Tasks” section lists a series of study activities that you should complete. Each chapter includes the activities that make the most sense for studying the topics in that chapter. The activities include the following:
 - **Review All Key Topics:** The Key Topic icon is shown next to the most important items in the “Foundation Topics” section of the chapter. The Review All Key Topics activity lists the key topics from the chapter and their corresponding page numbers. Although the contents of the entire chapter could be on the exam, you should definitely know the information listed in each key topic.
 - **Complete Tables and Lists from Memory:** To help you exercise your memory and memorize some facts, many of the more important lists and tables from the chapter are included in a document on the companion website. This document offers only partial information, allowing you to complete the table or list.
 - **Define Key Terms:** This section lists the most important terms from the chapter, asking you to write a short definition and compare your answer to the glossary at the end of this book.
- **Review Questions:** These questions at the end of each chapter measure insight into the topics that were discussed in the chapter.

- **End-of-Chapter Labs:** Real labs give you the right impression of what an exam assignment looks like. The end of chapter labs are your first step in finding out what the exam tasks really look like.

Other Features

In addition to the features in each of the core chapters, this book, as a whole, has additional study resources on the companion website, including the following:

- **Four practice exams:** *Red Hat RHCSA 8 Cert Guide* comes with four practice exams. You will find two in the book and two additional exams on the companion website; these are provided as PDFs so you can get extra practice testing your skills before taking the exam in the testing facility.
- **More than two hours of video training:** The companion website contains more than two hours of instruction from the best-selling *Red Hat Certified System Administrator (RHCSA) Complete Video Course*, 3rd Edition.

Exam Objective to Chapter Mapping

Table 1 details where every objective in the RHCSA exam is covered in this book so that you can more easily create a successful plan for passing the exam.

Table 1 Coverage of RHCSA Objectives

Objective	Chapter Title	Chapter
Access a shell prompt and issue commands with correct syntax	Using Essential Tools	2
Use input-output redirection (>, >>, , 2>, etc.)	Using Essential Tools	2
Use grep and regular expressions to analyze text	Working with Text Files	4
Access remote systems using SSH	Connecting to Red Hat Enterprise Linux 8	5
Log in and switch users in multiuser targets	Connecting to Red Hat Enterprise Linux 8	5
Archive, compress, unpack, and uncompress files using tar, star, gzip, and bzip2	Essential File Management Tools	3
Create and edit text files	Working with Text Files	4

Objective	Chapter Title	Chapter
Create, delete, copy, and move files and directories	Essential File Management Tools	3
Create hard and soft links	Essential File Management Tools	3
List, set, and change standard ugo/rwx permissions	Permissions Management	7
Locate, read, and use system documentation including man, info, and files in /usr/share/doc	Using Essential Tools	2
Boot, reboot, and shut down a system normally	Connecting to Red Hat Enterprise Linux 8	5
Boot systems into different targets manually	Essential Troubleshooting Skills	18
Interrupt the boot process in order to gain access to a system	Essential Troubleshooting Skills	18
Identify CPU/memory intensive processes and kill processes	Managing Processes	10
Adjust process scheduling	Managing Processes	10
Locate and interpret system log files and journals	Configuring Logging	13
Preserve system journals	Configuring Logging	13
Start, stop, and check the status of network services	Configuring Networking	8
Securely transfer files between systems	Connecting to Red Hat Enterprise Linux 8	5
List, create, and delete partitions on MBR and GPT disks	Managing Storage	14
Create and remove physical volumes	Managing Advanced Storage	15
Assign physical volumes to volume groups	Managing Advanced Storage	15
Create and delete logical volumes	Managing Advanced Storage	15
Configure systems to mount file systems at boot by universally unique ID (UUID) or label	Managing Storage	14
Add new partitions and logical volumes, and swap to a system non-destructively	Managing Storage	14
Create, mount, unmount, and use vfat, ext4, and xfs file systems	Managing Storage	14

Objective	Chapter Title	Chapter
Mount and unmount network file systems using NFS	Accessing Network Storage	24
Extend existing logical volumes	Managing Advanced Storage	15
Create and configure set-GID directories for collaboration	Permissions Management	7
Configure disk compression	Managing Advanced Storage	15
Manage layered storage	Managing Advanced Storage	15
Diagnose and correct file permission problems	Permissions Management	7
Schedule tasks using at and cron	Scheduling Tasks	12
Start and stop services and configure services to start automatically at boot	Working with Systemd	11
Configure systems to boot into a specific target automatically	Managing and Understanding the Boot Procedure	17
Configure time service clients	Configuring Time Services	25
Install and update software packages from Red Hat Network, a remote repository, or from the local file system	Managing Software	9
Work with package module streams	Managing Software	9
Modify the system bootloader	Managing and Understanding the Boot Procedure	17
Configure IPv4 and IPv6 addresses	Configuring Networking	8
Configure hostname resolution	Configuring Networking	8
Configure network services to start automatically at boot	Configuring Networking	8
Restrict network access using firewall-cmd/firewall	Configuring a Firewall	23
Create, delete, and modify local user accounts	User and Group Management	6
Change passwords and adjust password aging for local user accounts	User and Group Management	6
Create, delete, and modify local groups and group memberships	User and Group Management	6
Configure superuser access	User and Group Management	6

Objective	Chapter Title	Chapter
Manage security		
Configure firewall settings using firewall-cmd/firewalld	Configuring a Firewall	23
Configure key-based authentication for SSH	Configuring SSH	20
Set enforcing and permissive modes for SELinux	Managing SELinux	22
List and identify SELinux file and process context	Managing SELinux	22
Restore default file contexts	Managing SELinux	22
Use boolean settings to modify system SELinux settings	Managing SELinux	22
Diagnose and address routine SELinux policy violations	Managing SELinux	22

Where Are the Companion Content Files?

Register this print version of *Red Hat RHCSA 8 Cert Guide* to access the bonus content online.

This print version of this title comes with a website of companion content. You have online access to these files by following these steps:

1. Go to www.pearsonITcertification.com/register and log in or create a new account.
2. Enter the ISBN: **9780135938133**.
3. Answer the challenge question as proof of purchase.
4. Click the **Access Bonus Content** link in the Registered Products section of your account page to be taken to the page where your downloadable content is available.

Please note that many of the companion content files can be very large, especially image and video files.

If you are unable to locate the files for this title by following the steps, please visit www.pearsonITcertification.com/contact and select the Site Problems/Comments option. A customer service representative will assist you.

This book also includes an exclusive offer for 70% off the *Red Hat Certified System Administrator (RHCSA) Complete Video Course*, 3rd Edition.

Figure Credits

Chapter opener images by Charlie Edwards/ Photodisc/Getty Images

Figures 1-1 through 1-15 © 2019 Red Hat, Inc

Figures 2-1 and 2-2 © 2019 Red Hat, Inc

Figures 5-1 and 5-2 © 2019 Red Hat, Inc

Figures 8-1 through 8-5 © 2019 Red Hat, Inc

Figure 10-1 © 2019 Red Hat, Inc

Figures 18-1 through 18-5 © 2019 Red Hat, Inc

This page intentionally left blank



The following topics are covered in this chapter:

- Preparing to Install Red Hat Enterprise Linux
- Performing a Manual Installation

This chapter covers no exam objectives.

To learn how to work with Red Hat Enterprise Linux as an administrator, you first need to install it. This chapter teaches you how to set up an environment in which you can perform all exercises in this book.

On the RHCSA exam, you do not need to install Red Hat Enterprise Linux. However, because you need to install an environment that allows you to test all items discussed in this book, you start by installing Red Hat Enterprise Linux in this chapter. This chapter describes all steps that you will encounter while performing an installation of RHEL 8. It also discusses how to set up an environment in which you can perform all exercises in this book.

Because Red Hat Enterprise Linux contains registered trademarks, I will use CentOS to show you the installation. The only thing that is different in CentOS is the branding, which in CentOS is free to use. All software in CentOS is the same as the software in Red Hat Enterprise Linux. Because CentOS does not contain any registered trademarks, it is a common and popular solution that enables people to learn how to work with Red Hat Enterprise Linux.

Installing Red Hat Enterprise Linux

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 1-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 1-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Preparing to Install Red Hat Enterprise Linux	1, 2, 6
Performing a Manual Installation	3–5, 7–10

1. You want to install a test environment to practice for the RHCSA exam. Which of the following distributions should you avoid?
 - a. The most recent Fedora version
 - b. CentOS 8
 - c. Scientific Linux 8
 - d. RHEL 8
2. Which of the following features is available in both RHEL and CentOS?
 - a. Hardware certification
 - b. Software certification
 - c. The right to make support calls
 - d. Software updates

3. Why should you install the server with a GUI installation pattern?
 - a. To prepare for RHCSA, you need some tools that run in a GUI only.
 - b. The minimal installation is incomplete.
 - c. If you do not install a GUI immediately, it is hard to add it later.
 - d. The Server with GUI is the default installation that is recommended by Red Hat.
4. Which is the default file system that is used in RHEL 8?
 - a. Ext3
 - b. Ext4
 - c. XFS
 - d. Btrfs
5. Which feature is supported in Ext4 but not in XFS?
 - a. Shrinking the file system
 - b. Snapshots
 - c. File system quota
 - d. A maximum size that goes beyond 2 TB
6. Which of the following is not a reason why Fedora should be avoided?
 - a. Fedora contains features that may or may not be available in future RHEL releases.
 - b. Fedora distributions show a much later state of development than RHEL.
 - c. Fedora software is not stable.
 - d. Software in Fedora may differ from the same software in RHEL.
7. Which of the following options is not available from the Installation Summary screen?
 - a. Time & Date
 - b. Keyboard
 - c. Language Support
 - d. Troubleshoot an Existing Installation

8. After setting the root password that you want to use, you cannot proceed in the installation. What is the most likely reason?
 - a. The password is unsecure, and unsecure passwords are not accepted.
 - b. The password does not meet requirements in the password policy.
 - c. You also need to create a user.
 - d. If an unsecure password is used, you need to click Done twice.
9. Which statement about the system language is *not* true?
 - a. You can change the system language from the Installation Summary screen.
 - b. You can change the system language directly after booting from the installation media.
 - c. When setting the installation language, you can also select a keyboard layout.
 - d. After installation, you cannot change the language settings.
10. When installing a server that uses LVM logical volumes, you'll get at least three storage volumes (partitions or LVM). Which of the following is not part of them?
 - a. /boot
 - b. /var
 - c. /
 - d. swap

Foundation Topics

Preparing to Install Red Hat Enterprise Linux

Before you start installing Red Hat Enterprise Linux, a bit of preparation is helpful, as discussed in this section. You'll first learn what exactly Red Hat Enterprise Linux is. Then you'll learn how you can get access to the software. We then discuss the Red Hat Enterprise Linux add-ons, as well as the setup requirements. After you know all about these, you move on to the next section, where you learn how to install Red Hat Enterprise Linux.

What Is Red Hat Enterprise Linux 8 Server?

RHEL 8 is a Linux distribution. As you probably know, Linux is a free operating system. That means that the source code of all programs is available for free. However, some enterprise Linux distributions are sold as commercial products, with bundled support and maintenance, which is the case for RHEL 8; free downloads of RHEL 8 are restricted. There is a 60-day evaluation version available, and alternatively you can register for a free Red Hat developer subscription at <https://developers.redhat.com>.

To use RHEL 8, you need a subscription. This subscription entitles you to a few additional items, such as support and patches. When you pay for Red Hat Enterprise Linux, Red Hat offers you a supported Enterprise Linux operating system, which has some key benefits that are a normal requirement in corporate environments:

- Monitored updates and patches that have gone through a thorough testing procedure
- Different levels of support and help, depending on which type of subscription you have purchased
- A certified operating system that is guaranteed to run and to be supported on specific hardware models
- A certified platform for running enterprise applications such as SAP middle-ware, Oracle Database, and many more
- Access to the Red Hat Customer Portal at <https://access.redhat.com>, where you can find much detailed documentation that is available to customers only

Red Hat understands that not all potential customers are interested in these enterprise features. That is why Red Hat is involved in two free alternatives also:

- CentOS 8
- Fedora

You'll learn more about these in the upcoming sections of this chapter.

Getting the Software

There are different ways to get the software required to perform all exercises in this book. In this section, you learn what your options are.

Using Red Hat Enterprise Linux

If you want to learn how to work with the different programs, tools, and services that are provided in Red Hat Enterprise Linux 8, the easiest way is to use the developer program that Red Hat offers. Go to <https://developers.redhat.com> to register for the free developer program. This program gives you access to Red Hat Enterprise Linux for free, which allows you to work with RHEL in your own test environment without having to purchase it.

The most important things that you get in the official RHEL 8 Server release is access to the Red Hat Customer Portal. Through this portal, you have access to a wide variety of information regarding RHEL, in addition to updates provided through Red Hat Network (RHN). In particular, the Red Hat knowledge base is invaluable; you can use it to find answers to many common problems that have been posted there by Red Hat consultants.

Using CentOS

CentOS is the Community Enterprise Operating System. CentOS started as a recompiled version of RHEL, with all items that were not available for free removed from the RHEL software. Basically, just the name was changed and the Red Hat logo (which is proprietary) was removed from all the CentOS software packages. The result, still available today as CentOS 8, is an operating system that offers the same functionality as RHEL but is available for free (and without the enterprise support services).

In 2014, Red Hat incorporated CentOS, with the intent to offer something to customers who aren't ready for Enterprise support yet. The underlying idea is that eventually many customers will need enterprise support, because their Linux distribution will become increasingly important. And by giving away CentOS for free, and under the Red Hat brand, it is just natural that customers will upgrade to Red Hat Enterprise Linux.

CentOS 8 is also an excellent choice to work with in this book; it offers all that RHEL has to offer, but you do not have to pay for it. You can download CentOS 8 from <https://www.centos.org>. Note that CentOS doesn't always have the most recent version of the RHEL software and may run behind the latest RHEL release by between a couple of weeks and a couple of months.

Other Distributions

CentOS is not the only distribution that offers Red Hat packages without your having to pay for them. Another commonly used Linux distribution that was free—but based on the Red Hat Enterprise Linux packages—was Scientific Linux, a Linux distribution that was developed at Fermi National Accelerator Laboratory. Scientific Linux has been discontinued, and there is no Scientific Linux version that matches the Red Hat Enterprise Linux 8 software.

Another Linux distribution closely related to Red Hat Enterprise Linux is Fedora, a completely open source Linux distribution that is available for free. Red Hat has a lot of staff dedicated to contributing to the Fedora project, because Red Hat uses Fedora as the development platform for RHEL. The result is that Fedora offers access to the latest and greatest software, which in most cases is much more recent than the thoroughly tested software components of RHEL (which is why you should not use Fedora to prepare for the RHCSA exam). Fedora is also used by Red Hat as a testing ground for new features that might or might not be included in future RHEL releases. If you were to choose Fedora, you would be working with items that are not available in RHEL, which means that you would have to do things differently on the exam. So, don't use it!

Understanding Access to Repositories

An important difference between RHEL and the other distributions is the access to repositories. A *repository* is the installation source used for installing software. If you are using free software such as CentOS, correct repositories are automatically set up, and no further action is required. If you are using Red Hat Enterprise Linux with a subscription, you'll need to use the Subscription Manager software to get access to repositories.

TIP If you are installing Red Hat from the RHEL 8 installation disc, but you do not register it, you will not have access to a repository, which is why you need to know how to set up repository access manually. In Chapter 9, “Managing Software,” you learn how to do this.

Setup Requirements

RHEL 8 can be installed on physical hardware and on virtual hardware. For the availability of specific features, it does not really matter which type of hardware is used, as long as the following conditions are met:

- 1 GiB of RAM
- A 10-GiB hard disk
- A network card

TIP One GB is $1000 \times 1000 \times 1000$ bytes. With hardware vendors it is common to work with multiples of 1000; however, that doesn't correspond with how a computer works, which is why most computer software works with KiB, MiB, and GiB instead. In this, one GiB is $1024 \times 1024 \times 1024$ bytes (which is 1.07 GB).

The preceding requirements allow you to run a minimal installation of RHEL, but if you want to create an environment that enables you to perform all exercises described in this book, make sure to meet the following minimal requirements:

- 64-bit platform support
- 2 GiB of RAM
- A 20-GiB hard disk
- A DVD drive, either virtual or physical
- A network card

NOTE Some resources on the Internet will mention different minimal requirements. This is not a big deal for the RHCSA exam.

Cert Guide Environment Description

To set up an environment to work your way through this book, I suggest you start by installing one CentOS or RHEL 8 server, following the instructions in the next section. For the chapters in Part IV, “Managing Network Services,” it is useful if you have a second server as well. This second server doesn't have any specific requirements.

To set up the Cert Guide environment, I recommend that you use a solution for desktop virtualization, such as VMware Workstation (or Fusion if you are on Mac) or Oracle VirtualBox. Using one of these has the benefit that you can use snapshots,

which enables you to easily revert to a previous state of the configuration. Other virtualization solutions, such as KVM, are supported as well, but because KVM runs on Linux, you'll need to have some Linux knowledge already if you'd like to start with KVM. You can also install on real hardware, but that solution will be less flexible.

TIP In all chapters, you'll find step-by-step exercises that tell you exactly what to do to configure specific services. At the end of all chapters, you'll find end-of-chapter labs that provide assignments that are very similar to the types of assignments that you will encounter on the exam. To get the most out of the end-of-chapter labs, it is a good idea to start from a clean environment. The most efficient way to do this is by creating snapshots of the state of your virtual machines when you are starting the chapter. This allows you to revert to the state your virtual machines were in when you started working on the chapter, while still keeping all the work that you have done in previous chapters.

Performing a Manual Installation

Even if RHEL and CentOS 8 can be installed from other media such as an installation server or a USB key, the most common installation starts from the installation DVD or, when working in a virtual machine, from the installation DVD ISO file. So, take your installation DVD (or its ISO) and boot the computer on which you want to install the software. The following steps describe how to proceed from the moment you see the installation DVD boot screen. Notice that the steps describe how to install CentOS 8, and all branding in the figures is CentOS 8–based as well. The procedure does, however, apply to Red Hat Enterprise Linux 8 as well:

Key Topic

Step 1. After booting from DVD, you'll see the CentOS 8 boot menu. From this menu, you can choose from different options:

- **Install CentOS 8.0.0:** Choose this for a normal installation.
- **Test This Media & Install CentOS 8.0.0:** Select this if before installing you want to test the installation media. Note that this will take a significant amount of time.
- **Troubleshooting:** Select this option for some troubleshooting options. This option is useful if you cannot boot normally from your computer's hard drive.

When the installation program starts, you can pass boot options to the kernel to enable or disable specific features. To get access to the prompt where you can add these options, press Tab from the installation menu. This shows you the kernel boot line that will be used and offers an option to change boot parameters.

- Step 2.** To start a normal installation, select the **Install CentOS 8.0.0** boot option (see Figure 1-1). It will now load the installation system, and you'll be able to continue within a minute.




FIGURE 1-1 Select Install CentOS 8.0.0 to Start the Installation

- Step 3.** Once the base system from which you will perform the installation has loaded, you see the Welcome to CentOS 8.0.0 screen. From this screen, you can select the language and the keyboard setting. For the RHCSA exam, it makes no sense to choose anything but English. If you are working on a non-U.S. keyboard, from this screen you can select the keyboard setting. Make sure to select the appropriate keyboard setting, after which you click **Continue** to proceed (see Figure 1-2).
- Step 4.** After selecting the keyboard and language settings, you'll see the Installation Summary screen (see Figure 1-3). From this screen, you specify all settings you want to use. On this screen, you have several different options:
- **Keyboard:** Used to change the keyboard disposition.
 - **Language Support:** Used to add support for additional languages.
 - **Time & Date:** Used to specify the current time and date, as well as the time zone.
 - **Installation Source:** Used to specify where to install from. Typically, you'll install from the installation DVD.
 - **Software Selection:** Offers different installation patterns, to easily install a default set of packages.



FIGURE 1-2 Select the Appropriate Language and Keyboard Setting Before Continuing




FIGURE 1-3 Specify the Complete Configuration of Your Server from the Installation Summary Screen

- **Installation Destination:** Used to identify to which disk(s) to copy the files during the installation.
- **KDUMP:** Allows you to use a KDUMP kernel. This is a kernel that creates a core dump if anything goes wrong.
- **Network & Host Name:** Set IP address and related settings here.
- **Security Policy:** Offers a limited set of security policies, enabling you to easily harden a server.
- **System Purpose:** Allows you to indicate for which purpose this system is used. This information can easily be passed on to Red Hat support to facilitate any cases where support is needed.

From this Installation Summary screen, you can see whether items still need to be configured—these items are marked with an exclamation mark and a description in red text. As long as any issues exist, you cannot click the Begin Installation button (that is, it is disabled). You will not have to change settings for each option in all cases, but for completeness, the following steps describe the different settings available from the Installation Summary screen, with recommended changes where appropriate.

- Step 5.** Click the **Keyboard** option to view the settings to configure the keyboard layout. From this screen, you can also select a secondary keyboard layout, which is useful if your server is used by administrators using different keyboard layouts. Not only are different language settings supported, but also different hardware layouts. If many administrators are using an Apple Mac computer, for instance, you can select the standard keyboard layout for Mac in the appropriate region.

After adding another keyboard layout, you can also configure layout switching options. This is a key sequence that is used to switch between different kinds of layout. Select **Options** to specify the key combination you want to use for this purpose. After specifying the configuration you want to use, click **Done** to return to the Installation Summary screen.

- Step 6.** The Language Support option on the Installation Summary screen is the same as the Language Support option that you used in step 3 of this procedure. If you've already configured the language settings to be used, you do not need to change anything here.

Step 7. Click **Time & Date** to see a map of the world on which you can easily click the time zone that you are in (see Figure 1-4). Alternatively, you can select the region and city you are in from the corresponding drop-down list boxes. You can also set the current date and time, and after setting the network, you can specify the Network Time Protocol (NTP) to be used to synchronize time with time servers on the Internet. This option is not accessible if the network is not accessible. When using network time, you can add network time servers to be used by clicking the configuration icon in the upper-right part of the screen. After specifying the settings you want to use, click **Done** in the upper-left corner of the screen to write the settings.



FIGURE 1-4 Selecting Time & Date Settings

Step 8. In the Software section of the Installation Summary screen, click **Installation Source** to see the screen shown in Figure 1-5). If you have booted from a regular installation disc, there is nothing to specify. If you have booted from a minimal boot environment, you can specify the network URL where additional packages are available, as well as

additional repositories that need to be used. You do not have to do this for the RHCSA exam, but if ever you are setting up an installation server, it is useful to know that this option exists. Click **Done**.

The screenshot shows the 'INSTALLATION SOURCE' window for 'CENTOS LINUX 8.0.1905 INSTALLATION'. At the top left is a 'Done' button, and at the top right is a 'Help!' button. The main heading is 'Which installation source would you like to use?'. There are two radio buttons: 'Auto-detected installation media:' (selected) and 'On the network:'. Under 'Auto-detected installation media:', it shows 'Device: sr0' and 'Label: CentOS-8-BaseOS-x86_64' with a 'Verify' button. Under 'On the network:', there is a 'Closest mirror' dropdown, a text input field, a 'Proxy setup...' button, and a 'URL type:' dropdown. Below this is the 'Additional repositories' section. It contains a table with 'Enabled' and 'Name' columns. The 'AppStream' repository is checked and highlighted. To the right of the table are input fields for 'Name' (AppStream), 'file://' (/run/install/repo/AppStream), 'URL type:' (repository URL), 'Proxy URL:', 'User name:', and 'Password:'.

Enabled	Name
☒	AppStream

Additional repositories configuration fields:

- Name: AppStream
- file://: /run/install/repo/AppStream
- URL type: repository URL
- Proxy URL:
- User name:
- Password:

FIGURE 1-5 Selecting the Installation Source

Step 9. Click **Software Selection** to access an important part of the installation procedure (see Figure 1-6). From here, you select the base environment and choose add-ons that are available for the selected environment. The Minimal Install option is very common. This base environment allows you to install CentOS on a minimal-size hard disk. For this book, I assume that you install the server with the **Server with GUI** option. To perform the tasks that need to be performed on the RHCSA exam, some easy-to-use graphical tools are available, so it does make sense to install a server with a graphical user interface (GUI), even if you would never do this in a production environment. All additional packages can be added later. At this point, you do not have to select any additional packages. Click **Done**.




FIGURE 1-6 Make Sure You Select Server with GUI for Every Server You Are Going to Use for the Exercises in This Book

NOTE Some people say that *real* administrators do not use the Server with GUI installation pattern. Preparing for the RHCSA exam is not about being a real administrator. The big advantage of using the Server with GUI installation pattern is that it provides an easy-to-use interface. Some tools discussed in this book only run on a GUI. Also, when using a server with a GUI, you can use multiple terminal windows simultaneously, and that makes working with the RHEL command line really easy!

Step 10. After installing the software, you need to specify where you want to install to. Click **Installation Destination** on the Installation Summary screen. By default, automatic partitioning is selected, and you only need to approve the disk device you want to use for automatic partitioning (see Figure 1-7). Many advanced options are available, as well. To prepare your installation for all the exercises that are in later chapters in this book, you cannot just use the default partitioning. Instead, you need a setup that uses the Logical Volume Manager (LVM) and keeps some disk space available. To do this, from the screen you see in Figure 1-7, select the **Custom** radio button under Storage Configuration. Then, make sure that the disk you want to use is selected,

and click **Done** to proceed. Notice that sometimes not all options are shown, and you'll see a scrollbar to the right of the screen. If this is the case, scroll down to show additional installation options.




FIGURE 1-7 Select Custom and Click Done to Proceed

TIP If you want to use this server to perform the exercises in the upcoming chapters, make sure you keep some disk space that is not allocated to any partition. You need unpartitioned disk space to work through the partitioning and LVM exercises, which is an essential part of the RHCSA exam objectives.

After specifying that you want to set up disk layout manually, you'll see the Manual Partitioning screen shown in Figure 1-8. Click **+** to add new disk devices. To set up the environment that is required in this book, I recommend using the following disk layout (based on a 20-GiB hard disk):

- /boot mounted on an XFS-formatted traditional partition, size 500 MiB
- An XFS-formatted logical volume with a size of 10 GiB that is mounted on /
- A 1-GiB logical volume that is used as swap space




FIGURE 1-8 Configuring Advanced Disk Layout

NOTE RHEL 8 by default uses the XFS file system. This file system cannot be shrunk; it can only be expanded. Therefore, it is sometimes a better choice to use Ext4.

To create this configuration, after clicking **+** in the screen shown in Figure 1-8, you'll now see a pop-up dialog box in which you can specify a mount point and the desired capacity. From the Mount Point drop-down list, select **/boot** and add the desired capacity **500MiB**. Note that you can type **500M** instead, which will allocate a 500MiB capacity. Then, click **Add Mount Point**. You'll now go to the screen shown in Figure 1-9, in which you can specify specific details about the mount point you just created.

At this point, from the interface that you see in Figure 1-9, click the **+** sign again and specify the mount point **/** and a capacity of **10 GiB**. On the Mount Point Details screen that you see now, make sure that the device type LVM is set and that the file system XFS is selected. You do not have to modify anything else.




FIGURE 1-9 Specifying Details for the Mount Point You Just Created

Now click the **+** sign once more to add a swap device. From the Mount Point drop-down list, select **swap**, and specify the desired capacity of **1 GiB**. Then, click **Add Mount Point**. The layout at this point should look like Figure 1-10. If this is the case, click **Done** to write the configuration.

Before the configuration is committed to disk, you'll see a screen showing a summary of the changes that you have applied (see Figure 1-11). If it all looks good, click **Accept Changes** to write your configuration.




FIGURE 1-10 If the Configuration Looks Like This, You Can Write It to Disk




FIGURE 1-11 Before Writing the Changes to Disk, You Can Verify That It's All Good

Step 11. The next part of the Installation Summary screen enables you to set up networking. Notice that you must configure something. If you do not do anything, your server might not be able to connect to any network. Click **Network & Host Name** to set up networking. This opens the screen that you see in Figure 1-12.




FIGURE 1-12 On the Network & Host Name Screen, Set the Network Card to On

The network connection by default is set to Off. Make sure to switch it On. After switching on the network connection, you could click **Configure** to add further configuration. Networking is discussed in detail in Chapter 8, “Configuring Networking,” so you do not have to do that now and can just leave the default settings that get an IP address from the Dynamic Host Configuration Protocol (DHCP) server. You can also leave the hostname at its default setting; you learn how to change it in Chapter 8 as well.

Step 12. The Security Policy option enables you to tighten security on your server according to some common standards. Currently, the PCI-DSS and the OSPP options are supported. Open the interface and click **Select Profile** to apply either of these.

- Step 13.** The last option available in the Installation Summary enables you to define the System Purpose. Under this option you'll find some options that enable you to specify what you're using your server for. You can specify the role (such as RHEL Server or RHEL Workstation), as well as the Red Hat Service Level Agreement. Setting these options can make it easier to identify specific types of servers in a large corporate environment.
- Step 14.** After specifying all settings from the Installation Summary screen options, you can click **Begin Installation** to start the installation. This immediately starts the installation procedure but also prompts for user settings (see Figure 1-13).




FIGURE 1-13 Specifying User and Root Parameters

- Step 15.** From this screen, click **Root Password** first, and set the password to **password**. That is not very secure, but by using a simple password like this, you'll avoid issues later on that might result from not remembering

the password. You have to specify the password twice, and you also need to click **Done** twice (because you have to confirm that you really want to use a weak password).

- Step 16.** Next, click **User Creation** to create a user (see Figure 1-14). Enter the username **student**, and for this user, set the password to **password** also. Again, you have to click **Done** twice to confirm that you really want to use a weak password.

FIGURE 1-14 Specifying Additional User Settings

- Step 17.** When the installation has completed, you'll see the screen shown in Figure 1-15. You'll now need to click **Reboot** to restart the computer and finalize the installation.




FIGURE 1-15 Reboot to Finalize the Installation

Step 18. After rebooting, you have to go through a couple of additional setup steps. First, you need to accept the license agreement. To do this, click the red text **License Not Accepted**, select **I Accept the License Agreement**, and then click **Done** to complete. You can now click **Finish the Configuration** to finalize the configuration, which brings you to the graphical login prompt. Do NOT set up system registration using subscription-manager; we need a very specific setup and will take care of this in Chapter 9.

Summary

In this chapter, you learned what Red Hat Enterprise Linux is and how it relates to some other Linux distributions. You also learned how to install Red Hat Enterprise Linux 8 or CentOS 8. You are now ready to set up a basic environment that you can use to work on all the exercises in this book.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 1-2 lists a reference of these key topics and the page number on which each is found.

Table 1-2 Key Topics for Chapter 1

Key Topic Element	Description	Page
Step List	How to perform a manual RHEL 8 installation	10



Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

distribution, Linux, Red Hat, CentOS, Fedora, Scientific Linux

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. You do not want to buy an RHEL license, but you want to create an environment to practice for the exam. Which distribution should you use?
2. What happens to the network configuration if you don’t specify anything during the installation?
3. You want to install a minimal system. How much RAM do you need?
4. Why is it a good idea to have Internet access on all servers you are installing?
5. You want to install a virtual machine on a computer that does not have an optical disk drive. What is the easiest alternative to perform the installation?
6. Why is it a good idea to install a GUI?

7. What is the default file system on RHEL 8?
8. Can you install RHEL if you do not have Internet access?
9. What is the most important feature offered through RHN?
10. Which installation pattern should you use if you have a very limited amount of disk space available?

End-of-Chapter Lab

In this chapter, you learned how to set up Red Hat Enterprise Linux. At this point, you should have one server up and running. For exercises in later chapters in this book, one additional server is needed.

Lab 1.1

Repeat the procedure “Performing a Manual Installation” to install one more server. Details about the additional configuration on these servers follow in exercises in later chapters. For now, it is sufficient to ensure that the following conditions are met:

- Use the server names `server1` and `server2`.
- Set the network configuration to obtain an IP address automatically.
- Make sure to keep at least 1 GiB of disk space as unallocated disk space (which is not assigned to any partition) so that you have free space to work on the partitioning exercises in later chapters.
- Install one server using the Minimal Install pattern, and another server using the Server with GUI installation pattern.

This page intentionally left blank



The following topics are covered in this chapter:

- Basic Shell Skills
- Editing Files with vim
- Understanding the Shell Environment
- Finding Help

The following RHCSA exam objectives are covered in this chapter:

- Use input-output redirection (>, >>, |, 2>, etc.)
- Create and edit text files
- Locate, read, and use system documentation including man, info, and files in /usr/share/doc

This chapter is dedicated to coverage of the basic Linux skills that everyone should have before attempting to take the RHCSA exam.

Using Essential Tools

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 2-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 2-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Basic Shell Skills	1, 3, 4–7
Editing Files with vim	8
Understanding the Shell Environment	2, 9
Finding Help	10

- Which of the following commands enables you to redirect standard output as well as standard error to a file?
 - `1&2> file`
 - `> file 2>&1`
 - `>1&2 file`
 - `1>2& file`
- You are looking for a variable that is set in a Bash login shell for all users. Which of the following files is the most likely location where this variable is set? (Choose two.)
 - `/etc/profile`
 - `/etc/bashrc`
 - `~/bash_profile`
 - `~/bashrc`

3. A user has created a script with the name `myscript`. He tries to run it using the command **`myscript`**, but it is not started. The user has verified that the script permissions are set as executable. Which of the following is the most likely explanation?
 - a. An internal command is preventing the startup of the script.
 - b. Users are not allowed to run scripts.
 - c. The directory that contains the script is not in the `PATH`.
 - d. The script does not have appropriate permissions.
4. You need the output of the command `ls` to be used as input for the `less` command. Which of the following examples will do that for you?
 - a. `ls > less`
 - b. `ls >> less`
 - c. `ls >| less`
 - d. `ls | less`
5. A user wants to remove her complete history. Which of the following approaches would do that?
 - a. Remove the `~/.bash_history` file and type **`history -c`**.
 - b. Type **`history -c`**.
 - c. Remove the `~/.bash_history` file.
 - d. Type **`history -c`** and close the current shell.
6. Which of the following is *not* a valid method to repeat a command from history?
 - a. Press `Ctrl-r` and start typing a part of the command.
 - b. Type `!` followed by the first letters in the command.
 - c. Type `!` followed by the number of the command as listed in history.
 - d. Press `Ctrl-x` followed by the number in history.
7. For which of the following items can Bash completion be used?
 - a. Commands
 - b. Files
 - c. Variables
 - d. All of the above

8. Which of the following commands enables you to replace every occurrence of *old* with *new* in a text file that is opened with **vi**?
 - a. **:%s/old/new/g**
 - b. **:%r/old/new/**
 - c. **:s/old/new/g**
 - d. **r:/old/new**
9. Which approach works best if during the login process you want to show a message to all users who have just logged in to a shell session on your server?
 - a. Put the message in `/etc/issue`.
 - b. Put the message in `/etc/motd`.
 - c. Put the message in `/etc/profile`.
 - d. Put the message in `/etc/bashrc`.
10. You are using **man -k user**, but you get the message “nothing appropriate.” Which of the following solutions is most likely to fix this for you?
 - a. Type **updatedb** to update the mandb database.
 - b. Type **makewhatis** to update the mandb database.
 - c. Type **mandb** to update the mandb database.
 - d. Use **man -K**, not **man -k**.

Foundation Topics

Basic Shell Skills

The *shell* is the default working environment for a Linux administrator. It is the environment where users and administrators enter commands that are executed by the operating system. Different shells for Linux are available, but Bash is the most common shell. So when we are talking about “the shell” in this book, we are actually talking about the *Bash* shell. This chapter provides an overview of some of the items that you will encounter when working with the shell.

Understanding Commands

Working with the shell is all about working with command syntax. Typically, command syntax has three basic parts: the command, its options, and its arguments.

The command is the command itself, such as **ls**. This command shows a list of files in the current directory. To modify the behavior of the command, you can use *options*. Options are a part of the program code, and they modify what the command is doing. For instance, when you use the **-l** option with the **ls** command, a long listing of filenames and properties is displayed.

The word *argument* is a bit confusing. Generally speaking, it refers to anything that the command addresses, so anything you put after the command is an argument (including the options). Apart from the options that can be used as an argument, commands can have other arguments as well, which serve as a target to the command.

Let’s have a look at an example: the command **ls -l /etc**. This command has two different arguments: **-l** and **/etc**. The first argument is an option, modifying the behavior of the command. The second argument is a target, specifying where the command should do its work. You’ll find these three elements in nearly all commands you’ll be working with in a Linux environment.

Executing Commands

The purpose of the Linux shell is to provide an environment in which commands can be executed. The shell takes care of interpreting the command that a user has entered correctly. To do this, the shell makes a distinction between three kinds of commands:

- Aliases
- Internal commands
- External commands

An *alias* is a command that a user can define as needed. Some aliases are provided by default; type **alias** on the command line to get an overview. To define an alias, use **alias newcommand='oldcommand'**, as in the default alias **ll='ls -l --color=auto'** that has already been created on your system. Aliases are executed before anything else. So, if you have an alias with the name **ll** but also a command with the name **ll**, the alias will always take precedence, unless a complete pathname is used.

An *internal command* is a command that is a part of the shell itself and, as such, doesn't have to be loaded from disk separately. An *external command* is a command that exists as an executable file on the disk of the computer. Because it has to be read from disk the first time it is used, it is a bit slower. When a user executes a command, the shell first looks to determine whether it is an internal command; if it is not, it looks for an executable file with a name that matches the command on disk. To find out whether a command is a Bash internal or an executable file on disk, you can use the **type** command.

To look up external commands, use the **\$PATH** variable. This variable defines a list of directories that is searched for a matching filename when a user enters a command. To find out which exact command the shell will be using, you can use the **which** command. For instance, type **which ls** to find out where the shell will get the **ls** command from. An even stronger command is **type**. This command will also work on internal commands and aliases.

You should notice that for security reasons the current directory is not in the **\$PATH** variable and Linux does not look in the current directory to see whether a specific command is available from that directory. That is why you need to start a command that is in the current directory but nowhere in the **\$PATH** by including **./** in front of it. The dot stands for the current directory, and by running it as **./**, you tell Bash to look for the command in the current directory. Although running commands this way is not very common, you will have to do it to run scripts that you've created in your current directory.

The **\$PATH** variable can be set for specific users, but in general, most users will be using the same **\$PATH** variable. The only exception to this is the user root, who needs access to specific administration commands. In Exercise 2-1, you learn some of the basics about working with commands.

Exercise 2-1 Using Internal and External Commands from the Shell

1. Authenticate as the user who you created in Chapter 1, “Installing Red Hat Enterprise Linux” when installing your server.
2. Type **time ls**. This executes the **ls** command where the Bash internal **time** shows information about the time it took to complete this command.
3. Type **which time**. This shows the filename `/usr/bin/time` that was found in the **\$PATH** variable.
4. Type **echo \$PATH** to show the contents of the **\$PATH** variable. You can see that `/usr/bin` is included in the list, but because there also is an internal command **time**, the **time** command from the path will not be executed unless you tell the shell specifically to do so—the command in step 2 has executed the internal command for you because of command precedence.
5. Type **/usr/bin/time ls** to run the **/usr/bin/time** command when executing **ls**. You’ll notice that the output differs completely. Ignore the meaning of the output; we’ll get back to that later. What matters for now is that you realize that these are really two different commands.

I/O Redirection

By default, when a command is executed, it shows its results on the screen of the computer you are working on. The computer monitor is used as the standard destination for output, which is also referred to as the **STDOUT**. The shell also has default destinations to send error messages to and to accept input. Table 2-2 gives an overview of all three.

Table 2-2 Standard Input, Output, and Error Overview

Name	Default Destination	Use in Redirection	File Descriptor Number
STDIN	Computer keyboard	< (same as 0<)	0
STDOUT	Computer monitor	> (same as 1>)	1
STDERR	Computer monitor	2>	2

So if you run a command, that command would expect input from the keyboard, and it would normally send its output to the monitor of your computer without making a distinction between normal output and errors. Some commands, however, are started in the background and not from a current terminal session, so these commands do not have a monitor or console session to send their output to, and they do not listen to keyboard input to accept their standard input. That is where redirection comes in handy. Redirection is also useful if you want to work with input from an alternative location, such as a file.

Programs started from the command line have no idea what they are reading from or writing to. They just read from what the Linux kernel calls file descriptor 0 if they want to read from standard input, and they write to file descriptor number 1 to display non-error output (also known as “standard output”) and to file descriptor 2 if they have error messages to be output. By default, these file descriptors are connected to the keyboard and the screen. If you use redirection symbols such as `<`, `>`, and `|`, the shell connects the file descriptors to files or other commands. Let’s first look at `<` and `>`. Later we discuss pipes (the `|` symbol). Table 2-3 shows the most common redirectors that are used from the Bash shell.

Table 2-3 Common Bash Redirectors

Redirector	Explanation
<code>></code> (same as <code>1></code>)	Redirects STDOUT. If redirection is to a file, the current contents of that file are overwritten.
<code>>></code> (same as <code>1>></code>)	Redirects STDOUT. If output is written to a file, the output is appended to that file.
<code>2></code>	Redirects STDERR.
<code>2>&1</code>	Redirects STDERR to the same destination as STDOUT. Notice that this has to be used in combination with normal output redirection, as in <code>ls whuhui > errout 2>&1</code> .
<code><</code> (same as <code>0<</code>)	Redirects STDIN.

In I/O redirection, files can be used to replace the default STDIN, STDOUT, and STDERR. You can also redirect to *device files*. A device file on Linux is a file that is used to access specific hardware. Your hard disk, for instance, can be referred to as `/dev/sda`, the console of your server is known as `/dev/console` or `/dev/tty1`, and if you want to discard a command’s output, you can redirect to `/dev/null`. Note that to access most device files, you need to be root.

Using Pipes

Whereas an I/O redirector is used as an alternative for keyboard and computer monitor, a pipe can be used to catch the output of one command and use that as input for a second command. If a user runs the command **ls**, for instance, the output of the command is shown onscreen, because the screen is the default **STDOUT**. If the user uses **ls | less**, the commands **ls** and **less** are started in parallel. The standard output of the **ls** command is connected to the standard input of **less**. Everything that **ls** writes to the standard output will become available for read from standard input in **less**. The result is that the output of **ls** is shown in the **less** pager, where the user can browse up and down through the results easily.

As a Linux administrator, you'll use pipes a lot. Using pipes makes Linux a flexible operating system; by combining multiple commands using pipes, you can create "super" commands that make almost anything possible. In Exercise 2-2, you use I/O redirectors and pipes.

Exercise 2-2 Using I/O Redirection and Pipes

1. Open a shell as user **user** and type **cd** without any arguments. This ensures that the home directory of this user is the current directory while working on this exercise. Type **pwd** to verify this.
2. Type **ls**. You'll see the results onscreen.
3. Type **ls > /dev/null**. This redirects the **STDOUT** to the null device, with the result that you will not see it.
4. Type **ls ilwehgi > /dev/null**. This command shows a "no such file or directory" message onscreen. You see the message because it is not **STDOUT**, but rather an error message that is written to **STDERR**.
5. Type **ls ilwehgi 2> /dev/null**. Now you will no longer see the error message.
6. Type **ls ilwehgi Documents 2> /dev/null**. This shows the name of the Documents folder in your home directory while hiding the error message.
7. Type **ls ilwehgi Documents 2> /dev/null > output**. In this command, you still write the error message to **/dev/null** while sending **STDOUT** to a file with the name **output** that will be created in your home directory.
8. Type **cat output** to show the contents of this file.
9. Type **echo hello > output**. This overwrites the contents of the output file.
10. Type **ls >> output**. This appends the result of the **ls** command to the output file.

11. Type **ls -R /**. This shows a long list of files and folders scrolling over your computer monitor. (You might want to press Ctrl-C to stop [or wait some time]).
12. Type **ls -R | less**. This shows the same result, but in the pager **less**, where you can scroll up and down using the arrow keys on your keyboard.
13. Type **q** to close **less**. This will also end the **ls** program.
14. Type **ls > /dev/tty1**. This gives an error message because you are executing the command as an ordinary user and ordinary users cannot address device files directly (unless you were logged in to tty1). Only the user root has permission to write to device files directly.

History

A convenient feature of the Bash shell is the Bash history. Bash is configured by default to keep the last 1,000 commands a user used (and if the shell session is never closed, the exact number can grow well beyond that). When a shell session is closed, the history of that session is updated to the history file. The name of this file is `.bash_history`, and it is created in the home directory of the user who started a specific shell session. Notice that the history file is closed only when the shell session is closed; until that moment, all commands in the history are kept in memory.

The history feature makes it easy to repeat complex commands. There are several ways of working with history:

- Type **history** to show a list of all commands in the Bash history.
- Press Ctrl-r to open the prompt from which you can do backward searches in commands that you have previously used. Just type a string and Bash will look backward in the command history for any command containing that string as the command name or one of its arguments. Press Ctrl-r again to repeat the last backward search.
- Type **!*number*** to execute a command with a specific number from history.
- Type **!*sometext*** to execute the last command that starts with *sometext*. Notice that this is a potentially dangerous command because the command that was found is executed immediately!

In some cases it might be necessary to wipe the Bash history. This is useful, for instance, if you've typed a password in clear text by accident. If that happens, you can type **history -c** to clear the current history. Commands from this session won't be written to the history file when you exit the current session. If you want to remove the complete Bash history, type **history -w** immediately after using **history -c**.

Exercise 2-3 guides you through some history features.

Exercise 2-3 Working with History

1. Make sure that you have opened a shell as user **student**.
2. Type **history** to get an overview of commands that you have previously used.
3. Type some commands, such as the following:

```
ls
pwd
cat /etc/hosts
ls -l
```

The goal is to fill the history a bit.
4. Open a second terminal on your server by right-clicking the graphical desktop and selecting the **Open in Terminal** menu option.
5. Type **history** from this second terminal window. Notice that you do not see the commands that you just typed in the other terminal. That is because the history file has not been updated yet.
6. From the first terminal session, press Ctrl-r. From the prompt that opens now, type **ls**. You'll see the last **ls** command you used. Press Ctrl-r again. You'll now see that you are looking backward and that the previous **ls** command is highlighted. Press Enter to execute it.
7. Type **history | grep cat**. The **grep** command searches the history output for any commands that contain the text *cat*. Note the command number of one of the **cat** commands you have previously used.
8. Type **!nn**, where *nn* is replaced by the number you noted in step 7. You'll see that the last **cat** command is repeated.
9. Close this terminal by typing **exit**.
10. From the remaining terminal window, type **history -c**. This wipes all history that is currently in memory. Close this terminal session as well.
11. Open a new terminal session and type **history**. It may be a bit unexpected, but you'll see a list of commands anyway. That is because **history -c** clears the in-memory history, but it does not remove the `.bash_history` file in your home directory.

Bash Completion

Another useful feature of the Bash shell is command-line completion. This feature helps you in finding the command you need, and it also works on variables and file-names. If you have installed the **bash-completion** software package, it works for some of the more complex commands as well.

Bash completion is useful when working with commands. Just type the beginning of a command and press the Tab key on your computer's keyboard. If there is only one option for completion, Bash will complete the command automatically for you. If there are several options, you need to press Tab once more to get an overview of all the available options. In Exercise 2-4, you learn how to work with these great features.

Exercise 2-4 Using Bash Completion

1. Still from a user shell, type **gd** and press Tab. You'll see that nothing happens.
2. Press Tab again. Bash now shows a short list of all commands that start with the letters *gd*.
3. To make it clear to Bash what you want, type **i** (so that your prompt at this point shows the command **gdi**). Press Tab again. Bash now knows what you want and opens **gdisk** for you. Press Enter to close the prompt that was just opened.
4. Use **cd /etc** to go to the /etc directory.
5. Type **cat pas** and press Tab. Because there is one file only that starts with *pas*, Bash knows what to do and automatically completes the filename. Press Enter to execute the command.

Editing Files with vim

Managing Linux often means working with files. Most things that are configured on Linux are configured through files. To complete administrative tasks, you often need to change the contents of a configuration file with a text editor.

Over the years, many text editors have been created for Linux. One editor really matters, though, and that is **vi**. Even if some other text editors are easier to use, **vi** is the only text editor that is always available. That is why as a Linux administrator you need to know how to work with **vi**. One common alternative is **vim**, or “**vi** improved”; it is a complete rewrite of **vi** with a lot of enhancements that make working with **vi** easier, such as syntax highlighting for many configuration files, which makes it easy to recognize typing errors that you have made. Everything that you learn in this section about **vim** works on **vi** as well.

An important concept when working with **vim** is that it uses different modes. Two of them are particularly important: *command mode* and *input mode*. These modes often cause confusion because in command mode you can just enter a command and you cannot change the contents of a text file. To change the contents of a text file, you need to get to input mode.

The challenge when working with **vim** is the vast number of commands that are available. Some people have even produced **vim** cheat sheets, listing all available commands. Do not use them. Instead, focus on the minimal number of commands that are really important. Table 2-4 summarizes the most essential **vim** commands. Use these (and only these) and you'll do fine on the RHCSA exam.

TIP Do *not* try to work with as many commands as possible when working with **vim**. Just use a minimal set of commands and use them often. You'll see; you'll get used to these commands and remember them on the exam. Also, you may like the **vimtutor** command. (Use **yum install vim-enhanced** to install it.) This command opens a **vim** tutorial that has you work through some nice additional exercises.

Key
Topic

Table 2-4 vim Essential Commands

vim Command	Explanation
Esc	Switches from input mode to command mode. Press this key before typing any command.
i, a	Switches from command mode to input mode at (i) or after (a) the current cursor position.
o	Opens a new line below the current cursor position and goes to input mode.
:wq	Writes the current file and quits.
:q!	Quits the file without applying any changes. The ! forces the command to do its work. Only add the ! if you really know what you are doing.
:w filename	Writes the current file with a new filename.
dd	Deletes the current line.
yy	Copies the current line.
p	Pastes the current selection.
v	Enters visual mode, which allows you to select a block of text using the arrow keys. Use d to cut the selection or y to copy it.
u	Undoes the last command. Repeat as often as necessary.
Ctrl-r	Redoes the last undo.
gg	Goes to the first line in the document.
G	Goes to the last line in the document.

vim Command	Explanation
/text	Searches for <i>text</i> from the current cursor position forward.
?text	Searches for <i>text</i> from the current cursor position backward.
^	Goes to the first position in the current line.
\$	Goes to the last position in the current line.
!ls	Adds the output of ls (or any other command) in the current file.
:%s/old/new/g	Replaces all occurrences of <i>old</i> with <i>new</i> .

Note that you know the most essential commands for working with **vim**. Exercise 2-5 gives you the opportunity to test them.

Exercise 2-5 vim Practice

1. Type **vim ~/testfile**. This starts **vim** and opens a file with the name `testfile` in `~`, which represents your current home directory.
2. Type **i** to enter input mode and then type the following text:


```
cow
sheep
ox
chicken
snake
fish
oxygen
```
3. Press **Esc** to get back to command mode and type **:w** to write the file using the same filename.
4. Type **:3** to go to line number 3.
5. Type **dd** to delete this line.
6. Type **dd** again to delete another line.
7. Type **u** to undo the last deletion.
8. Type **o** to open a new line.
9. Enter some more text at the current cursor position:


```
tree
farm
```
10. Press **Esc** to get back into command mode.
11. Type **:%s/ox/OX/g**.
12. Type **:wq** to write the file and quit. If for some reason that does not work, use **:wq!**.

Understanding the Shell Environment

When you are working from a shell, an environment is created to ensure that all that is happening is happening the right way. This environment consists of variables that define the user environment, such as the **\$PATH** variable discussed earlier. In this section, you get a brief overview of the shell environment and how it is created.

Understanding Variables

The Linux shell environment consists of many variables. *Variables* are fixed names that can be assigned dynamic values. An example of a variable is **\$LANG**, which in my shell is set to **en_US.UTF-8**. This value (which may differ on your system) ensures that I can work in the English language using settings that are common in the English language (think of how date and time are displayed).

The advantage for scripts and programs of working with variables is that the program only has to use the name of the variable without taking interest in the specific value that is assigned to the variable. Because different users have different needs, the variables that are set in a user environment will differ. To get an overview of the current variables defined in your shell environment, type the **env** command, which will show environment variables. Example 2-1 shows some lines of the output of this command.

Example 2-1 Displaying the Current Environment

```
[user@server1 ~]$ env
MAIL=/var/spool/mail/user
PATH=/usr/local/bin:/bin:/usr/bin:/usr/local/sbin:/usr/sbin:/home/
user/.local/bin:/home/user/bin
PWD=/home/user
LANG=en_US.UTF-8
HISTCONTROL=ignoredups
SHLVL=1
HOME=/home/user
LOGNAME=user
LESSOPEN=||/usr/bin/lesspipe.sh %s
_=/bin/env
OLDPWD=/etc
```

As you can see from Example 2-1, to define a variable, you type the name of the variable, followed by an equal sign (=) and the value that is assigned to the specific variable. To read the value of a variable, you can use the **echo** command (among others), followed by the name of the variable, as in **echo \$PATH**, which reads the current value of the **\$PATH** variable and prints that on the STDOUT. For now, you do not have to know much more about variables. You can read about more advanced use of variables in Chapter 19, “An Introduction to Bash Shell Scripting.”

Recognizing Environment Configuration Files

When a user logs in, an environment is created for that user automatically. This happens based on four different configuration files where some script code can be specified and where variables can be defined for use by one specific user:

- **/etc/profile:** This is the generic file that is processed by all users upon login.
- **/etc/bashrc:** This file is processed when subshells are started.
- **~/.bash_profile:** In this file, user-specific login shell variables can be defined.
- **~/.bashrc:** In this user-specific file, subshell variables can be defined.

As you have seen, in these files a distinction is made between a login shell and a subshell. A login shell is the first shell that is opened for a user after the user has logged in. From the login shell, a user may run scripts, which will start a subshell of that login shell. Bash allows for the creation of a different environment in the login shell and in the subshell, but to make sure the same settings are used in all shells, it's a good idea to include subshell settings in the login shell as well.

Using /etc/motd and /etc/issue

To display messages during the login process, Bash uses the **/etc/motd** and the **/etc/issue** files. Messages in **/etc/motd** display after a user has successfully logged in to a shell. (Note that users in a graphical environment do not see its contents after a graphical login.) Using **/etc/motd** can be a convenient way for system administrators to inform users about an issue or a security policy, for example.

Another way to send information to users is by using **/etc/issue**. The contents of this file display before the user logs in. This provides an excellent means of specifying specific login instructions to users who are not logged in yet.

In Exercise 2-6, you can review the topics that have been discussed in this section.

Exercise 2-6 Managing the Shell Environment

1. Open a shell in which you are user **user**.
2. Type **echo \$LANG** to show the contents of the variable that sets your system keyboard and language settings.
3. Type **ls --help**. You'll see that help about the **ls** command is displayed in the current language settings of your computer.
4. Type **LANG=es_ES.UTF-8**. This temporarily sets the language variable to Spanish.
5. Type **ls --help** again. You'll see that now the **ls** help text is displayed in Spanish.
6. Type **exit** to close your terminal window. Because you have not changed the contents of any of the previously mentioned files, when you open a new shell, the original value of the **LANG** variable will be used.
7. Open a shell as **user** again.
8. Verify the current value of the **LANG** variable by typing **echo \$LANG**.
9. Type **vim .bashrc** to open the **.bashrc** configuration file.
10. In this file, add the line **COLOR=red** to set a variable with the name **COLOR** and assign it the value **red**. Notice that this variable doesn't really change anything on your system; it just sets a variable.
11. Close the user shell and open a new user shell.
12. Verify that the variable **COLOR** has been set, by using **echo \$COLOR**. Because the **.bashrc** file is included in the login procedure, the variable is set after logging in.

Finding Help

On an average Linux system, hundreds of commands are available—way too many to ever be able to remember all of them, which is why using the help resources on your computer is so very important. The **man** command is the most important resource for getting help about command syntax and usage. Apart from that, you can show a compact list of command options using **command --help**.

Using --help

The quickest way to get an overview of how to use a command is by running the command with the **--help** option. Nearly all commands will display a usage summary when using this option. The list of options that is shown in this way is of use mainly when you already have a generic understanding of how to use the command

and need a quick overview of options available with the command—it doesn’t give detailed information that will help users who don’t know the command yet.

TIP Nearly all commands provide a short overview of help when the option **--help** is used. Some commands do not honor that option and consider it erroneous. Fortunately, these commands will be so friendly as to show an error message, displaying valid options with the command, which effectively means that you’ll get what you needed anyway.

Using man

When using the Linux command line, you will at some point consult man pages. The **man** command is what makes working from the command line doable. If you do not know how a command is used, the man page of that command will provide valuable insight. This section covers a few **man** essentials.

To start with, the most important parts of the man page in general are at the bottom of the man page. Here you’ll find two important sections: In many cases there are examples; if there are no examples, there is always a “See Also” section. The topics you find here are related man pages, which is useful if you have just not hit the right man page. To get to the bottom of the man page as fast as possible, use the **G** command. You can also type **/example** to search the man page for any examples. Figure 2-1 shows what the end of a man page may look like.

```

File Edit View Search Terminal Help
Create a 1TiB thin LV, first creating a new thin pool for it, where the thin pool has 100MiB of space, uses 2 stripes, has a
64KiB stripe size, and 256KiB chunk size.
lvcreate --type thin --name mylv --thinpool mypool
-V 1t -L 100m -i 2 -I 64 -c 256 vg00

Create a thin snapshot of a thin LV (the size option must not be used, otherwise a copy-on-write snapshot would be created).
lvcreate --snapshot --name mysnap vg00/thinvol

Create a thin snapshot of the read-only inactive LV named "origin" which becomes an external origin for the thin snapshot LV.
lvcreate --snapshot --name mysnap --thinpool mypool vg00/origin

Create a cache pool from a fast physical device. The cache pool can then be used to cache an LV.
lvcreate --type cache-pool -L 1G -n my_cpool vg00 /dev/fast1

Create a cache LV, first creating a new origin LV on a slow physical device, then combining the new origin LV with an existing
cache pool.
lvcreate --type cache --cachepool my_cpool
-L 100G -n mylv vg00 /dev/slow1

SEE ALSO
lvm(8) lvm.conf(5) lvmconfig(8)

pvchange(8) pvck(8) pvcreate(8) pvdisplay(8) pvmove(8) pvremove(8) pvresize(8) pvs(8) pvscan(8)

vgcfgbackup(8) vgcfgrestore(8) vgchange(8) vgck(8) vgcreate(8) vgconvert(8) vgdisplay(8) vgexport(8) vgextend(8) vgimport(8)
vgimportclone(8) vgmerge(8) vgmknodes(8) vgreduce(8) vgrename(8) vgs(8) vgscan(8) vgsplit(8)

lvcreate(8) lvchange(8) lvconvert(8) lvdisplay(8) lvextend(8) lvreduce(8) lvremove(8) lvrename(8) lvresize(8) lvs(8) lvscan(8)

lvm-fullreport(8) lvm-lvpoll(8) lvm2-activation-generator(8) blkdeactivate(8) lvmddump(8)

dmeventd(8) lvmcmd(8) lvmlockd(8) lvmlockctl(8) clvmd(8) cmirror(8) lvmdbusd(8)

lvmsystemid(7) lvmreport(7) lvraid(7) lvthin(7) lvmcache(7)

Red Hat, Inc. LVM TOOLS 2.03.00(2) (2018-10-10) LVMCREATE(8)
Manual page lvcreate(8) line 1910/1945 (END) (press h for help or q to quit)

```

FIGURE 2-1 Sample man Page Contents

Finding the Right man Page

To find information in man pages, you can search the mandb database by using **apropos** or **man -k**. If the database is current, getting access to the information you need is easy. Just type **man -k**, followed by the keyword you want to search for. This command looks in the summary of all man pages that are stored in the mandb database. Example 2-2 shows a partial result of this command.

Example 2-2 Searching man Pages with **man -k**

```
[root@server1 ~]# man -k partition
addpart (8)      - simple wrapper around the "add partition" ioctl
cfdisk (8)      - display or manipulate disk partition table
cgdisk (8)      - Curses-based GUID partition table (GPT)
                  manipulator
delpart (8)     - simple wrapper around the "del partition" ioctl
fdisk (8)       - manipulate disk partition table
fixparts (8)    - MBR partition table repair utility
gdisk (8)       - Interactive GUID partition table (GPT) manipulator
iostat (1)      - Report Central Processing Unit (CPU) statistics
                  and in...
kpartx (8)      - Create device maps from partition tables
mpartition (1)  - partition an MSDOS hard disk
os-prober (1)   - Discover bootable partitions on the local system
partprobe (8)   - inform the OS of partition table changes
partx (8)       - tell the Linux kernel about the presence and
                  numbering...
pvcreate (8)    - initialize a disk or partition for use by LVM
pvresize (8)    - resize a disk or partition in use by LVM2
resizepart (8)  - simple wrapper around the "resize partition" ioctl
sfdisk (8)      - partition table manipulator for Linux
sgdisk         (- Command-line GUID partition table (GPT)
                  manipulator fo...
systemd-efi-boot-generator (8) - Generator for automatically mounting
                  the EFI...
systemd-gpt-auto-generator (8) - Generator for automatically
                  discovering and ..
```

Based on the information that **man -k** is giving you, you can probably identify the man page that you need to access to do whatever you want to accomplish. Be aware, however, that **man -k** is not perfect; it searches only the short summary of each command that is installed. If your keyword is not in the summary, you'll find nothing and get a "nothing appropriate" error message.

TIP Instead of using **man -k**, you can use the **apropos** command, which is equivalent to **man -k**.

When using **man -k** to find specific information from the man pages, you'll sometimes get a load of information. If that happens, it might help to filter down the results a bit by using the **grep** command. But if you want to do that, it is important that you know what you are looking for.

Man pages are categorized in different sections. The most relevant sections for system administrators are as follows:

- **1**: Executable programs or shell commands
- **5**: File formats and conventions
- **8**: System administration commands

Key
Topic

There are also sections that provide in-depth details about your Linux system, such as the sections about system calls and library calls. When using **man -k**, you'll get results from all of these sections. To limit the results that display, it makes sense to use **grep** to show only those sections that are relevant for what you need. So, if you are looking for the configuration file that has something to do with passwords, use **man -k password | grep 5**, or if you are looking for the command that an administrator would use to create partitions, use **man -k partition | grep 8**.

Another useful man option is **-f**. The command **man -f <somecommand>** displays a short description of the item as found in the mandb database. This may help you when deciding whether this man page contains the information you are looking for.

Updating mandb

As previously mentioned, when using the **man -k** command, the mandb database is consulted. This database is automatically created through a **cron** scheduled job. Occasionally, you might look for something that should obviously be documented but all you get is the message “nothing appropriate.” If that happens, you might need to update the mandb database. Doing that is easy: Just run the **mandb** command as root without any arguments. It will see whether new man pages have been installed and update the mandb database accordingly.

TIP Do not try to memorize all the commands that you need to accomplish specific tasks. Instead, memorize how to find these commands and find which man page to read to get more information about the command. In Exercise 2-7, you'll see how that works.

Assume that you are looking for a command, using **man -k**, but all you get is the message “nothing appropriate” and you do not remember how to fix it. Exercise 2-7 shows what you can do in such cases.

Exercise 2-7 Using man -k

1. Because **man -k** does not give the expected result, it makes sense to look in the man page for the **man** command for additional information about **man -k**. Type **man man** to open the man page of **man**. Once in the man page, type **/-k** to look for a description of the **-k** option. Type **n** a few times until you get to the line that describes the option. You’ll see that **man -k** is equivalent to **apropos** and that you can read the man page of **apropos** for more details. So type **q** to exit this man page.
2. Type **man apropos** and read the first paragraphs of the description. You’ll see that the database searched by **apropos** is updated by the **mandb** program.
3. Type **man mandb**. This man page explains how to run **mandb** to update the mandb database. As you’ll read, all you need to do is type **mandb**, which does the work for you.
4. Type **mandb** to update the mandb database. Notice that you won’t see many man pages being added if the mandb database was already quite accurate.

Using info

Apart from the information that you’ll find in man pages, another system provides help about command usage. This is the info system. Most commands are documented in man pages, but some commands have their main documentation in the info system and only show a short usage summary in the man page. If that is the case, the “See Also” section of the man page of that command will tell you that “The full documentation for...is maintained as a Texinfo manual.” You then can read the Info page using the command **pinfo** or **info**. Both commands work, but in **pinfo**, special items such as menu items are clearly indicated, which is why using **pinfo** is easier.

When working with **info**, take a look at the top line of the viewer. This shows the current position in the info document. Particularly interesting are the Up, Next, and Previous indicators, which tell you how to navigate. Info pages are organized like web pages, which means that they are organized in a hierarchical way. To browse through that hierarchy, type **n** to go to the next page, **p** to go to the previous page, or **u** to move up in the hierarchy.

In an info page, you’ll also find menus. Each item that is marked with an asterisk (*) is a menu item. Use the arrow keys to select a specific menu item. This brings you down one level. To get back up again, type **u**. This brings you back to the original starting point in the **pinfo** hierarchy. Figure 2-2 shows what an info page looks like.

```

File: coreutils.info, Node: Introduction, Next: Common options, Prev: Top, Up: Top
1 Introduction
*****

This manual is a work in progress: many sections make no attempt to
explain basic concepts in a way suitable for novices. Thus, if you are
interested, please get involved in improving this manual. The entire
GNU community will benefit.

The GNU utilities documented here are mostly compatible with the
POSIX standard.

Please report bugs to <bug-coreutils@gnu.org>. Include the version
number, machine architecture, input files, and any other information
needed to reproduce the bug: your input, what you expected, what you
got, and why it is wrong.

If you have a problem with 'sort' or 'date', try using the '--debug'
option, as it can often help find and fix problems without having to
wait for an answer to a bug report. If the debug output does not
suffice to fix the problem on your own, please compress and attach it to
the rest of your bug report.

Although diffs are welcome, please include a description of the
problem as well, since this is sometimes difficult to infer. *Note
(gcc)@bugs::.

This manual was originally derived from the Unix man pages in the
distributions, which were written by David MacKenzie and updated by Jim
Meyering. What you are reading now is the authoritative documentation
for these utilities; the man pages are no longer being maintained. The
original 'fmt' man page was written by Ross Paterson. François Pinard
did the initial conversion to Texinfo format. Karl Berry did the
indexing, some reorganization, and editing of the results. Brian
Yomans of the Free Software Foundation office staff combined the
Viewing line 36/40, 90%

```

FIGURE 2-2 Getting More Command Usage Information Using **pinfo**

Exercise 2-8 shows an example of such a command, and in this exercise you learn how to get the information out of the info page.

Exercise 2-8 Using info

1. Type **man ls**. Type **G** to go to the end of the man page and look at the “See Also” section. It tells you that the full documentation for **ls** is maintained as a Texinfo manual. Quit the man page by pressing **q**.
2. Type **pinfo coreutils 'ls invocation'**. This shows the information about **ls** usage in the **pinfo** page. Read through it and press **q** when done.

Using /usr/share/doc Documentation Files

A third source of information consists of files that are sometimes copied to the /usr/share/doc directory. This happens in particular for services and larger systems that are a bit more complicated. You will not typically find much information about a command like **ls**, but some services do provide useful information in /usr/share/doc.

Some services store very useful information in this directory, like rsyslog, bind, Kerberos, and OpenSSL. For some services, even example files are included. One example of these services is VDO, which is covered in more detail in Chapter 15, “Managing Advanced Storage.”

Summary

In this chapter, you read about essential Linux administration tasks. You learned about some of the important shell basics, such as I/O redirection, working with history, and management of the environment. You also learned how to edit text files with the **vim** editor. In the last part of this chapter, you learned how to find information using **man** and related commands.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 2-5 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 2-5 Key Topics for Chapter 2

Key Topic Element	Description	Page
Table 2-4	vim essential commands	40
List	Significant sections in man	47

Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

shell, bash, internal command, external command, **\$PATH**, variable, STDIN, STDOUT, STDERR, file descriptor, pipe, redirect, device files, environment, login shell, subshell

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. What is a variable?
2. Which command enables you to find the correct man page based on keyword usage?
3. Which file do you need to change if you want a variable to be set for every shell that is started?
4. When analyzing how to use a command, you read that the documentation is maintained with the Techinfo system. How can you read the information?
5. What is the name of the file where Bash stores its history?
6. Which command enables you to update the database that contains man keywords?
7. How can you undo the last modification you have applied in **vim**?
8. What can you add to a command to make sure that it does not show an error message, assuming that you do not care about the information that is in the error messages either?
9. How do you read the current contents of the **\$PATH** variable?
10. How do you repeat the last command you used that contains the string *dog* somewhere in the command?

End-of-Chapter Lab

You have now learned about some of the most important basic skills that a Linux administrator should have. In this section, you apply these skills by doing an end-of-chapter lab.

Lab 2.1

1. Modify your shell environment so that on every subshell that is started, a variable is set. The name of the variable should be **COLOR**, and the value should be set to **red**. Verify that it is working.
2. Use the appropriate tools to find the command that you can use to set the system time 1 minute ahead.
3. From your home directory, type the command **ls -al werghl *** and ensure that errors as well as regular output are redirected to a file with the name `/tmp/lsoutput`.



This chapter covers the following subjects:

- Working with the File System Hierarchy
- Managing Files
- Using Links
- Working with Archives and Compressed Files

The following RHCSA exam objectives are covered in this chapter:

- Create, delete, copy, and move files and directories
- Archive, compress, unpack, and uncompress files using **tar**, **star**, **gzip**, and **bzip2**
- Create hard and soft links

Linux is a file-oriented operating system. That means that many things an administrator has to do on Linux can be traced down to managing files on the Linux operating system. Also, when using hardware devices, files are involved. This chapter introduces you to essential file management skills. You learn how the Linux file system is organized and how you can work with files and directories. You also learn how to manage links and compressed or uncompressed archives.

Essential File Management Tools

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 3-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 3-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Working with the File System Hierarchy	1–4
Managing Files	5–7
Using Links	8
Working with Archives and Compressed Files	10

- Under which directory would you expect to find nonessential program files?
 - /boot
 - /bin
 - /sbin
 - /usr
- Under which directory would you expect to find log files?
 - /proc
 - /run
 - /var
 - /usr

3. Which of the following directories would typically not be mounted on its own dedicated device?
 - a. /etc
 - b. /boot
 - c. /home
 - d. /usr
4. Which of the following commands would give the most accurate overview of mounted disk devices (without showing much information about mounted system devices as well)?
 - a. **mount**
 - b. **mount -a**
 - c. **df -hT**
 - d. **du -h**
5. Which command enables you to show all files in the current directory so that the newest files are listed last?
 - a. **ls -lRt**
 - b. **ls -lrt**
 - c. **ls -alrt**
 - d. **ls -alr**
6. Which command enables you to copy hidden files as well as regular files from /home/\$USER to the current directory?
 - a. **cp -a /home/\$USER .**
 - b. **cp -a /home/\$USER/* .**
 - c. **cp -a /home/\$USER/. .**
 - d. **cp -a home/\$USER. .**
7. Which command enables you to rename the file myfile to mynewfile?
 - a. **mv myfile mynewfile**
 - b. **rm myfile mynewfile**
 - c. **rn myfile mynewfile**
 - d. **ren myfile mynewfile**

8. Which statement about hard links is *not* true?
 - a. Hard links cannot be created to directories.
 - b. Hard links cannot refer to files on other devices.
 - c. The inode keeps a hard link counter.
 - d. If the original hard link is removed, all other hard links become invalid.
9. Which command creates a symbolic link to the directory /home in the directory /tmp?
 - a. `ln /tmp /home`
 - b. `ln /home /tmp`
 - c. `ln -s /home /tmp`
 - d. `ln -s /tmp /home`
10. Which `tar` option enables you to add one single file to a tar archive?
 - a. `-a`
 - b. `-A`
 - c. `-r`
 - d. `-u`

Foundation Topics

Working with the File System Hierarchy

To manage a Linux system, you should be familiar with the default directories that exist on almost all Linux systems. This section describes these directories and explains how mounts are used to compose the file system hierarchy.

Defining the File System Hierarchy

The file system on most Linux systems is organized in a similar way. The layout of the Linux file system is defined in the Filesystem Hierarchy Standard (FHS), and this file system hierarchy is described in **man 7 hier**. Table 3-2 shows an overview of the most significant directories that you'll encounter on a Red Hat Enterprise Linux (RHEL) system, as specified by the FHS.



Table 3-2 FHS Overview

Directory	Use
/	The root directory. This is where the file system tree starts.
/boot	Contains all files and directories that are needed to boot the Linux kernel.
/dev	Contains device files that are used for accessing physical devices. This directory is essential during boot.
/etc	Contains configuration files that are used by programs and services on your server. This directory is essential during boot.
/home	Used for local user home directories.
/media, /mnt	Contain directories that are used for mounting devices in the file system tree.
/opt	Used for optional packages that may be installed on your server.
/proc	Used by the proc file system. This is a file system structure that gives access to kernel information.
/root	The home directory of the root user.
/run	Contains process and user-specific information that has been created since the last boot.
/srv	May be used for data by services like NFS, FTP, and HTTP.
/sys	Used as an interface to different hardware devices that is managed by the Linux kernel and associated processes.

Directory	Use
/tmp	Contains temporary files that may be deleted without any warning during boot.
/usr	Contains subdirectories with program files, libraries for these program files, and documentation about them.
/var	Contains files that may change in size dynamically, such as log files, mail boxes, and spool files.

Understanding Mounts

To understand the organization of the Linux file system, the concept of mounting is important. A Linux file system is presented as one hierarchy, with the root directory (/) as its starting point. This hierarchy may be distributed over different devices and even computer systems that are mounted into the root directory.

In the process of mounting, a device is connected to a specific directory, such that after a successful mount this directory gives access to the device contents.

Mounting devices makes it possible to organize the Linux file system in a flexible way. There are several disadvantages to storing all files in just one file system, which gives several good reasons to work with multiple mounts:

- High activity in one area may fill up the entire file system, which will negatively impact services running on the server.
- If all files are on the same device, it is difficult to secure access and distinguish between different areas of the file system with different security needs. By mounting a separate file system, mount options can be added to meet specific security needs.
- If a one-device file system is completely filled, it may be difficult to make additional storage space available.

To avoid these pitfalls, it is common to organize Linux file systems in different devices (and even shares on other computer systems), such as disk partitions and logical volumes, and mount these devices into the file system hierarchy. By configuring a device as a dedicated mount, it is also possible to use specific mount options that can restrict access to the device. Some directories are commonly mounted on dedicated devices:

- **/boot:** This directory is often mounted on a separate device because it requires essential information your computer needs to boot. Because the root directory (/) is often on a Logical Volume Manager (LVM) logical volume, from which Linux cannot boot by default, the kernel and associated files need to be stored separately on a dedicated /boot device.

- **/var:** This directory is often on a dedicated device because it grows in a dynamic and uncontrolled way (for example, because of the log files that are written to **/var/log**). By putting it on a dedicated device, you can ensure that it will not fill up all storage on your server.
- **/home:** This directory often is on a dedicated device for security reasons. By putting it on a dedicated device, it can be mounted with specific options, such as **noexec** and **nodev**, to enhance the security of the server. When reinstalling the operating system, it is an advantage to have home directories in a separate file system. The home directories can then survive the system reinstall.
- **/usr:** This directory contains operating system files only, to which normal users normally do not need any write access. Putting this directory on a dedicated device allows administrators to configure it as a read-only mount.

Apart from these directories, you may find servers that have other directories that are mounted on dedicated partitions or volumes also. After all, it is up to the discretion of the administrator to decide which directories get their own dedicated devices.

To get an overview of all devices and their mount points, you can use different commands:

- The **mount** command gives an overview of all mounted devices. To get this information, the **/proc/mounts** file is read, where the kernel keeps information about all current mounts. It shows kernel interfaces also, which may lead to a long list of mounted devices being displayed. Example 3-1 shows sample output of this command.

Example 3-1 Partial **mount** Command Output

```
[root@server1 ~]# mount
sysfs on /sys type sysfs (rw,nosuid,nodev,noexec,relatime,seclabel)
proc on /proc type proc (rw,nosuid,nodev,noexec,relatime)
devtmpfs on /dev type devtmpfs (rw,nosuid,seclabel,size=909060k,
nr_inodes=227265,mode=755)
securityfs on /sys/kernel/security type securityfs (rw,nosuid,nodev,
noexec,relatime)
tmpfs on /dev/shm type tmpfs (rw,nosuid,nodev,seclabel)
devpts on /dev/pts type devpts (rw,nosuid,noexec,relatime,seclabel,
gid=5,mode=620,ptmxmode=000)
tmpfs on /run type tmpfs (rw,nosuid,nodev,seclabel,mode=755)
tmpfs on /sys/fs/cgroup type tmpfs (ro,nosuid,nodev,noexec,seclabel,
mode=755)
...
```

```

/dev/nvme0n1p1 on /boot type xfs (rw,relatime,seclabel,attr2,inode64,
noquota)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw,relatime)
tmpfs on /run/user/42 type tmpfs (rw,nosuid,nodev,relatime,seclabel,
size=184968k,mode=700,uid=42,gid=42)
tmpfs on /run/user/1000 type tmpfs (rw,nosuid,nodev,relatime,seclabel,
size=184968k,mode=700,uid=1000,gid=1000)
gvfsd-fuse on /run/user/1000/gvfs type fuse.gvfsd-fuse
(rw,nosuid,nodev,relatime,user_id=1000,group_id=1000)
/dev/sr0 on /run/media/student/RHEL-8-0-BaseOS-x86_64 type iso9660
(ro,nosuid,nodev,relatime,nojoliet,check=s,map=n,blocksize=2048,
uid=1000,gid=1000,dmode=500,fmode=400,uhelper=udisks2)
tmpfs on /run/user/0 type tmpfs (rw,nosuid,nodev,relatime,seclabel,
size=184968k,mode=700))

```

- The **df -Th** command was designed to show available disk space on mounted devices; it includes most of the system mounts. Because it will look on all mounted file systems, it is a convenient command to get an overview of current system mounts. The **-h** option summarizes the output of the command in a human-readable way, and the **-T** option shows which file system type is used on the different mounts.
- The **findmnt** command shows mounts and the relationship that exists between the different mounts. Because the output of the **mount** command is a bit overwhelming, you may like the output of **findmnt**. Example 3-2 shows sample output of this command. Notice that because of width limitations of the book page, the output that belongs in the **OPTIONS** column appears on the left side of the page.

Example 3-2 Sample **findmnt** Command Output

```

[root@server1 ~]# findmnt

```

TARGET	SOURCE	FSTYPE	OPTIONS
/	/dev/mapper/rhel-root		
		xfs	
rw,relatime,seclabel,at			
└-/sys	sysfs	sysfs	
rw,nosuid,nodev,noexec,			
└-/sys/kernel/security	securityfs	securit	
rw,nosuid,nodev,noexec,			
└-/sys/fs/cgroup	tmpfs	tmpfs	
ro,nosuid,nodev,noexec,			
...			

└─/proc	proc	proc
rw,nosuid,nodev,noexec,		
└─/proc/sys/fs/binfmt_misc	systemd-1	autofs
rw,relatime,fd=46,pgrp=		
└─/dev	devtmpfs	devtmpfs
rw,nosuid,seclabel,size		
└─/dev/shm	tmpfs	tmpfs
rw,nosuid,nodev,seclabe		
└─/dev/pts	devpts	devpts
rw,nosuid,noexec,relati		
└─/dev/mqueue	mqueue	mqueue
rw,relatime,seclabel		
└─/dev/hugepages	hugetlbfs	hugetlb
rw,relatime,seclabel,pa		
└─/run	tmpfs	tmpfs
rw,nosuid,nodev,seclabe		
└─/run/user/0	tmpfs	tmpfs
rw,nosuid,nodev,relatim		
└─/run/user/42	tmpfs	tmpfs
rw,nosuid,nodev,relatim		
└─/run/user/1000	tmpfs	tmpfs
rw,nosuid,nodev,relatim		
└─└─/run/user/1000/gvfs	gvfsd-fuse	fuse.gv
rw,nosuid,nodev,relatim		
└─/run/media/student/RHEL-8-0-BaseOS-x86_64		
	/dev/sr0	iso9660
ro,nosuid,nodev,relatim		
└─/boot	/dev/nvme0n1p1	
		xfs
rw,relatime,seclabel,at		
└─/var/lib/nfs/rpc_pipefs	sunrpc	rpc_pipe rw,relatime

In Exercise 3-1, you use different commands to get an overview of currently mounted devices.

Exercise 3-1 Getting an Overview of Current Mounts

1. Log in as an ordinary user and type **mount**. Notice that the output of the command is quite overwhelming. If you read carefully, though, you'll see a few directories from the Linux directory structure and their corresponding mounts.
2. Now type **df -hT**. Notice that a lot fewer devices are shown. An example of the output of this command is shown in Example 3-3.

Example 3-3 `df -hT` Sample Output

```
[root@server1 ~]# df -hT
```

Filesystem	Type	Size	Used	Avail	Use%	Mounted on
/dev/mapper/centos-root	xfs	5.9G	3.9G	2.1G	66%	/
devtmpfs	devtmpfs	908M	0	908M	0%	/dev
tmpfs	tmpfs	918M	144K	917M	1%	/dev/shm
tmpfs	tmpfs	918M	21M	897M	3%	/run
tmpfs	tmpfs	918M	0	918M	0%	/sys/fs/cgroup
/dev/sda1	xfs	197M	131M	67M	67%	/boot

Now that you have entered the **mount** and **df** commands, let's have a closer look at the output of the **df -hT** command in Example 3-3.

The output of **df** is shown in seven columns:

- **Filesystem:** The name of the device file that interacts with the disk device that is used. The real devices in the output start with `/dev` (which refers to the directory that is used to store device files). You can also see a couple of `tmpfs` devices. These are kernel devices that are used to create a temporary file system in RAM.
- **Type:** The type of file system that was used.
- **Size:** The size of the mounted device.
- **Used:** The amount of disk space the device has in use.
- **Avail:** The amount of unused disk space.
- **Use%:** The percentage of the device that currently is in use.
- **Mounted on:** The directory the device currently is mounted on.

Note that when using the **df** command, the sizes are reported in kibibytes. The option **-m** will display these in mebibytes, and using **-h** will display a human-readable format in KiB, MiB, GiB, TiB, or PiB.

Managing Files

As an administrator, you need to be able to perform common file management tasks. These tasks include the following:

- Working with wildcards
- Managing and working with directories
- Working with absolute and relative pathnames

- Listing files and directories
- Copying files and directories
- Moving files and directories
- Deleting files and directories

The following subsections explain how to perform these tasks.

Working with Wildcards

When working with files, using wildcards can make your work a lot easier. A wildcard is a shell feature that helps you refer to multiple files in an easy way. Table 3-3 gives an overview.

Key
Topic

Table 3-3 Wildcard Overview

Wildcard	Use
*	Refers to an unlimited number of any characters. ls * , for instance, shows all files in the current directory (except those that have a name starting with a dot).
?	Used to refer to one specific character that can be any character. ls c?t would match <i>cat</i> as well as <i>cut</i> .
[auo]	Refers to one character that may be selected from the range that is specified between square brackets. ls c[auo]t would match <i>cat</i> , <i>cut</i> , and <i>cot</i> .

Managing and Working with Directories

To organize files, Linux works with directories (also referred to as folders). You have already read about some default directories as defined by the FHS. When users start creating files and storing them on a server, it makes sense to provide a directory structure as well. As an administrator, you have to be able to walk through the directory structure. Exercise 3-2 gives you practice working with directories.

Exercise 3-2 Working with Directories

1. Open a shell as a normal user. Type **cd**. Next, type **pwd**, which stands for *print working directory*. You'll see that you are currently in your home directory; that is, name `/home/<username>`.
2. Type **touch file1**. This command creates an empty file with the name `file1` on your server. Because you currently are in your home directory, you can create any file you want to.

3. Type **cd /**. This changes the current directory to the root (/) directory. Type **touch file2**. You'll see a "permission denied" message. Ordinary users can create files only in directories where they have the permissions needed for this.
4. Type **cd /tmp**. This brings you to the /tmp directory, where all users have write permissions. Again, type **touch file2**. You'll see that you can create items in the /tmp directory (unless there is already a file2 that is owned by somebody else).
5. Type **cd** without any arguments. This command brings you back to your home directory.
6. Type **mkdir files**. This creates a directory with the name files in the current directory. The **mkdir** command uses the name of the directory that needs to be created as a relative pathname; it is relative to the position you are currently in.
7. Type **mkdir /home/\$USER/files**. In this command, you are using the variable **\$USER**, which is substituted with your current username. The complete argument of **mkdir** is an absolute filename to the directory files you are trying to create. Because this directory already exists, you'll get a "file exists" error message.
8. Type **rmdir files** to remove the directory files you have just created. The **rmdir** command enables you to remove directories, but it works only if the directory is empty and does not contain any files.

Working with Absolute and Relative Pathnames

In the previous section, you worked with the commands **cd** and **mkdir**. You used these commands to browse through the directory structure. You also worked with a relative filename and an absolute filename.

An *absolute filename*, or *absolute pathname*, is a complete path reference to the file or directory you want to work with. This pathname starts with the root directory, followed by all subdirectories up to the actual filename. No matter what your current directory is, absolute filenames will always work. An example of an absolute filename is /home/lisa/file1.

Key
Topic

A *relative filename* is relative to the current directory as shown with the **pwd** command. It contains only the elements that are required to get from the current directory up to the item you need. Suppose that your current directory is /home (as shown by the **pwd** command). When you refer to the relative filename lisa/file1, you are referring to the absolute filename /home/lisa/file1.

Key
Topic

When working with relative filenames, it is sometimes useful to move up one level in the hierarchy. Imagine you are logged in as root and you want to copy the file /home/lisa/file1 to the directory /home/lara. A few solutions would work:

- Use **cp /home/lisa/file1 /home/lara**. Because in this command you are using absolute pathnames, this command will work at all times.

- Make sure your current directory is `/home` and use **`cp lisa/file1 lara`**. Notice that both the source file and the destination file are referred to as relative filenames and for that reason do *not* start with a `/`.
- If the current directory is set to `/home/lisa`, you could also use **`cp file1 ../lara`**. In this command, the name of the target file uses `..`, which means go up one level. The `..` is followed by `/lara`, so the total name of the target file would be interpreted as “go up one level” (so you would be in `/home`), and from there, look for the `/lara` subdirectory.

TIP If you are new to working with Linux, understanding relative filename is not always easy. There is an easy workaround, though. Just make sure that you always work with absolute pathnames. It is more typing, but it is easier so you’ll make fewer mistakes.

In Chapter 2, “Using Essential Tools,” you learned how you can use Bash completion via the Tab key to complete commands. Using Bash completion makes it a lot easier to work with long commands. Bash completion works on filenames, too. If you have a long filename, like `my-long-file-name`, try typing **`my-`** and press the Tab key. If in the current directory, just one file has a name starting with `my-`, the filename will automatically be completed. If there are more files that have a name starting with `my-`, you have to press the Tab key twice to see a list of all available filenames.

Listing Files and Directories

While working with files and directories, it is useful to show the contents of the current directory. For this purpose, you can use the **`ls`** command. If used without arguments, **`ls`** shows the contents of the current directory. Some common arguments make working with **`ls`** easier. Table 3-4 gives an overview.

Key Topic

Table 3-4 **`ls`** Common Command-Line Options

Command	Use
<code>ls -l</code>	Shows a long listing, which includes information about file properties, such as creation date and permissions.
<code>ls -a</code>	Shows all files, including hidden files.
<code>ls -lrt</code>	This is a very useful command. It shows commands sorted on modification date. You’ll see the most recently modified files last in the list.
<code>ls -d</code>	Shows the names of directories, not the contents of all directories that match the wildcards that have been used with the <code>ls</code> command.
<code>ls -R</code>	Shows the contents of the current directory, in addition to all of its subdirectories; that is, it R ecursively descends all subdirectories.

TIP A hidden file on Linux is a file that has a name that starts with a dot. Try the following: **touch .hidden**. Next, type **ls**. You will not see the file. Then type **ls -a**. You'll see it.

When using **ls** and **ls -l**, you'll see that files are color-coded. The different colors that are used for different file types make it easier to distinguish between different kinds of files. Do not focus too much on them, though, because the colors that are used are the result of a variable setting that might be different in other Linux shells or on other Linux servers.

Copying Files

To organize files on your server, you'll often copy files. The **cp** command helps you do so. Copying a single file is not difficult: Just use **cp /path/to/file /path/to/destination**. To copy the file `/etc/hosts` to the directory `/tmp`, for instance, use **cp /etc/hosts /tmp**. This results in the file `hosts` being written to `/tmp`.

TIP If you copy a file to a directory, but the target directory does not exist, a file will be created with the name of the alleged target directory. In many cases, that's not the best solution and it would be better to just get an error message instead. You can accomplish this by placing a **/** after the directory name, so use **cp /etc/hosts /tmp/** and not **cp /etc/hosts /tmp**.

With the **cp** command, you can also copy an entire subdirectory, with its contents and everything beneath it. To do so, use the option **-R**, which stands for recursive. (You'll see the option **-R** with many other Linux commands also.) For example, to copy the directory `/etc` and everything in it to the directory `/tmp`, you would use the command **cp -R /etc /sound**.

While using the **cp** command, you need to consider permissions and other properties of the files. Without extra options, you risk these properties not being copied. If you want to make sure that you keep the current permissions, use the **-a** option, which has **cp** work in archive mode. This option ensures that permissions and all other file properties will be kept while copying. So, to copy an exact state of your home directory and everything within it to the `/tmp` directory, use **cp -a ~ /tmp**.

A special case when working with **cp** is hidden files. By default, hidden files are not copied over. There are three solutions to copy hidden files as well:

- **cp /somedir/.*/tmp** This copies all files that have a name starting with a dot (the hidden files, that is) to `/tmp`. It gives an error message for directories whose name starts with a dot in `/somedir`, because the **-R** option was not used.

- **cp -a /somedir/ .** This copies the entire directory /somedir, including its contents, to the current directory. So, as a result, a subdirectory somedir will be created in the current directory.
- **cp -a /somedir/. .** This copies all files, regular and hidden, to the current directory (notice the space between the two dots at the end of this command).

Moving Files

To move files, you use the **mv** command. This command removes the file from its current location and puts it in the new location. You can also use it to rename a file (which, in fact, is nothing else than copying and deleting the original file anyway). Let's take a look at some examples:

- **mv myfile /tmp:** Moves the file myfile from the current directory to /tmp.
- **mkdir somefiles; mv somefiles /tmp:** First creates a directory with the name somefiles and then moves this directory to /tmp. Notice that this also works if the directory contains files.
- **mv myfile mynewfile:** Renames the file myfile to a new file with the name mynewfile.

Deleting Files

The last common file administration task is file deletion. To delete files and directories, you use the **rm** command. When used on a single file, the single file is removed. You can also use it on directories that contain files. To do so, include the **-r** option, which again stands for recursive.

NOTE Many commands have an option that creates recursive behavior. On some commands you use the option **-R**, and on other commands you use the option **-r**. That is confusing, but it is just the way it is.

On RHEL 8, if you use the **rm** command as root, it prompts for confirmation. This is because through /root/.bashrc, **rm** is defined as an alias to **rm -i**. If you do not like that, you can use the **-f** option or remove the alias from /root/.bashrc. Make sure that you know what you are doing when using this option, because after using it, there is no way back but the backup tape!

In Exercise 3-3, you work with the common file management utilities.

NOTE In this exercise dots are important and used as a part of the commands. To avoid confusion, if normally a dot would be used to indicate the end of a sentence, in this exercise we've left it out if it immediately follows a command.

Exercise 3-3 Working with Files

Figure 3-1 provides an overview of the directory structure you are working with in this exercise.



FIGURE 3-1 Sample Directory Structure Overview

1. Open a shell as an ordinary user.
2. Type **pwd**
You should be in the directory `/home/$USER`.
3. Type **mkdir newfiles oldfiles**
Type **ls**
You'll see the two directories you have just created.
4. Type **touch newfiles/.hidden** and **touch newfiles/unhidden**
This creates two files in the directory `newfiles`.
5. Type **cd oldfiles**
6. Type **ls -al**
This shows two items only: `.`, which refers to the current directory; and `..`, which refers to the item above this (the parent directory).
7. Type **ls -al ../newfiles**
In this command, you are using a relative pathname to refer to the contents of the `/home/$USER/newfiles` directory.
8. Use the command **cp -a ../newfiles/ .** (notice the space between the `/` and the `.` at the end of the command).
9. Type **ls -a**
You see that you have created the subdirectory `newfiles` into the directory `oldfiles`.
10. Make sure that you are still in `/home/$USER/oldfiles`, and type **rm -rf newfiles**

11. Now use the command **cp -a ../newfiles/* ..** (notice the space between the ***** and **..**). Type **ls -al** to see what has been copied now. You'll see that the hidden file has not been copied.
12. To make sure that you copy hidden files as well as regular files, use **cp -a ../newfiles/ ..**
13. Verify the command worked this time, using **ls -al**
You'll notice that the hidden files as well as the regular files have been successfully copied.

Using Links

Links on Linux are like aliases that are assigned to a file. There are symbolic links, and there are hard links. To understand a link, you need to know a bit about how the Linux file system uses inodes for file system administration.

Understanding Hard Links

Key Topic

Linux stores administrative data about files in *inodes*. The inode is used to store all administrative data about files. Every file on Linux has an inode, and in the inode, important information about the file is stored:

- The data block where the file contents are stored
- The creation, access, and modification date
- Permissions
- File owners

Just one important piece of information is not stored in the inode: the name of the file. Names are stored in the directory, and each filename knows which inode it has to address to access further file information. It is interesting to know that an inode does not know which name it has; it just knows how many names are associated with the inode. These names are referred to as *hard links*.

When you create a file, you give it a name. Basically, this name is a hard link. On a Linux file system, multiple hard links can be created to a file. This can be useful, because it enables you to access the file from multiple different locations. Some restrictions apply to hard links, though:

- Hard links must exist all on the same device (partition, logical volume, etc).
- You cannot create hard links to directories.
- When the last name (hard link) to a file is removed, access to the file's data is also removed.

The nice thing about hard links is that no difference exists between the first hard link and the second hard link. They are both just hard links, and if the first hard link that ever existed for a file is removed, that does not impact the other hard links that still exist. The Linux operating system uses links on many locations to make files more accessible.

Understanding Symbolic Links

A *symbolic link* (also referred to as *soft link*) does not link directly to the inode but to the name of the file. This makes symbolic links much more flexible, but it also has some disadvantages. The advantage of symbolic links is that they can link to files on other devices, as well as on directories. The major disadvantage is that when the original file is removed, the symbolic link becomes invalid and does not work any longer.

Figure 3-2 gives a schematic overview of how inodes, hard links, and symbolic links relate to one another.




FIGURE 3-2 Links and Inodes Overview

Creating Links

Use the **ln** command to create links. It uses the same order of parameters as **cp** and **mv**; first you mention the source name, followed by the destination name. If you want to create a symbolic link, you use the option **-s**, and then you specify the source and target file or directory. One important restriction applies, however: to be able to create hard links, you must be the owner of the item that you want to link to. This is a relatively new security restriction that was introduced in RHEL 7.

Table 3-5 shows some examples.

Table 3-5 ln Usage Examples	
Command	Explanation
<code>ln /etc/hosts .</code>	Creates a link to the file /etc/hosts in the current directory
<code>ln -s /etc/hosts .</code>	Creates a symbolic link to the file /etc/hosts in the current directory
<code>ln -s /home /tmp</code>	Creates a symbolic link to the directory /home in the directory /tmp

Key
Topic

The **ls** command will reveal whether a file is a link:

- In the output of the **ls -l** command, the first character is an **l** if the file is a symbolic link.
- If a file is a symbolic link, the output of **ls -l** shows the name of the item it links to after the filename.
- If a file is a hard link, **ls -l** shows the hard link counter. In the output in Example 3-4, this is the number 3 that is right before `root root` for the `hosts` file.

Example 3-4 Showing Link Properties with **ls -l**

```
[root@localhost tmp]# \ls -l
total 3
lrwxrwxrwx. 1 root root 5 Jan 19 04:38 home -> /home
-rw-r--r--. 3 root root 158 Jun 7 2013 hosts
```

NOTE In Example 3-4, the command used is **\ls -l**, not **ls -l**. The **ls** command by default is an alias, which takes care of using the different colors when showing **ls** output; the **** in front of the command causes the alias not to be used.

Removing Links

Removing links can be dangerous. To show you why, let's consider the following procedure.

1. Make a directory named `test` in your home directory: **mkdir ~/test**.
2. Copy all files that have a name starting with `a`, `b`, `c`, `d`, or `e` from `/etc` to this directory: **cp /etc/[a-e]* ~/test**.
3. Make sure that you are in your home directory, by using **cd** without arguments.
4. Type **ln -s test link**.
5. Type **rm link**. This removes the link. (Do *not* use **-r** or **-f** to remove links, even if they are subdirectories.)
6. Type **ls -l**. You'll see that the symbolic link has been removed.
7. Let's do it again. Type **ln -s test link** to create the link again.

8. Type **rm -rf link/** (which is what you would get by using Bash command-line completion).
9. Type **ls**. You'll see that the directory link still exists.
10. Type **ls test/**. You'll see the directory test is now empty.

In Exercise 3-4, you learn how to work with symbolic links and hard links.

Exercise 3-4 Working with Symbolic Links and Hard Links

1. Open a shell as a regular (nonroot) user.
2. From your home directory, type **ln /etc/passwd ..** (Make sure that the command ends with a dot!) This command gives you an “operation not permitted” error because you are not the owner of /etc/passwd.
3. Type **ln -s /etc/passwd ..** (Again, make sure that the command ends with a dot!) This works; you do not have to be the owner to create a symbolic link.
4. Type **ln -s /etc/hosts** (this time with no dot at the end of the command). You'll notice this command also works. If the target is not specified, the link is created in the current directory.
5. Type **touch newfile** and create a hard link to this file by using **ln newfile linkedfile**.
6. Type **ls -l** and notice the link counter for newfile and linkedfile, which is currently set to 2.
7. Type **ln -s newfile symlinkfile** to create a symbolic link to newfile.
8. Type **rm newfile**.
9. Type **cat symlinkfile**. You will get a “no such file or directory” error message because the original file could not be found.
10. Type **cat linkedfile**. This gives no problem.
11. Type **ls -l** and look at the way the symlinkfile is displayed. Also look at linkedfile, which now has the link counter set to 1.
12. Type **ln linkedfile newfile**.
13. Type **ls -l** again. You'll see that the original situation has been restored.

Working with Archives and Compressed Files

Another important file-related task is managing archives and compressed files. To create an archive of files on a Linux computer, the **tar** command is often used. This

command was originally designed to stream files to a tape without any compression of the files. If you want to compress files as well, a specific compression tool has to be used, or you need to specify an option that compresses the archive while it is created. In this section, you learn how to work with archives and compressed files.

Managing Archives with **tar**

The Tape ARchiver (**tar**) utility is used to archive files. Although originally designed to stream files to a backup tape, in its current use **tar** is used mostly to write files to an archive file. You have to be able to perform four important tasks with **tar** on the RHCSA exam:

- Create an archive
- List the contents of an archive
- Extract an archive
- Compress and uncompress archives

Creating Archives with **tar**

To create an archive, you use the **tar -cf archivename.tar /files-you-want-to-archive** command. If you want to see what is happening, use the **-v** option as well. To put files in an archive, you need at least read permissions to the file and execute permissions on the directory the file resides in. Use **tar -cvf /root/homes.tar /home** as user root to write the contents of the /home directory and everything below it to the file homes.tar in the directory /root. Notice the options that are used; the order in these options is important.

Originally, **tar** did not use the dash (-) in front of its options. Modern **tar** implementations use that dash, as do all other Linux programs, but they still allow the old usage without a dash for backward compatibility. For a complete overview of relevant options used, see Table 3-6.

While managing archives with **tar**, it is also possible to add a file to an existing archive or to update an archive. To add a file to an archive, you use the **-r** options. Use, for instance, **tar -rvf /root/homes.tar /etc/hosts** to add the /etc/hosts file to the archive.

To update a currently existing archive file, you can use the **-u** option. So, use **tar -uvf /root/homes.tar /home** to write newer versions of all files in /home to the archive.

Monitoring and Extracting tar Files

Before extracting a file, it is good to know what might be expected. The option **-t** can be used to find out. Type, for instance, **tar -tvf /root/homes.tar** to see the contents of the tar archive.

TIP It is good practice to create archive files with an extension such as **.tar** or **.tgz** so that they can be easily recognized, but not everyone does that. If you think that a file is a tar archive, but you are not sure, use the **file** command. If you type **file somefile**, for instance, the **file** command analyzes its contents and shows on the command line what type of file it is.

To extract the contents of an archive, use **tar -xvf /archivename**. This extracts the archive in the *current* directory. That means that if you are in **/root** when typing **tar -xvf /root/homes.tar**, and the file contains a directory **/home**, after extracting you'll have a new directory **/root/home** that contains the entire contents of the file. This might not be what you wanted to accomplish. There are two solutions to put the extracted contents right where you want to have them:

- Before extracting the archive file, **cd** to the directory where you want to extract the file.
- Use the option **-C /targetdir** to specify the target directory where you want to extract the file to. If you want to put the contents of the file **/root/homes.tar** in the directory **/tmp**, for instance, you can use **tar -xvf homes.tar -C /tmp**.

NOTE The RHCSA objectives mention that you need to know how to work with **star** as well. The **star** utility was designed to offer support for archiving nondefault file attributes, such as access control lists (see Chapter 7, “Permissions Management”) or SELinux file context (see Chapter 22, “Managing SELinux”). In its current release, **tar** offers this functionality also, so there is no real need to use **star** anymore.

Apart from extracting an entire archive file, it is also possible to extract one file out of the archive. To do so, use **tar -xvf /archivename.tar file-you-want-to-extract**. If your archive **etc.tar** contains the file **/etc/hosts** that you want to extract, for instance, use **tar -xvf /root/etc.tar etc/hosts**.

Using Compression

Many files contain a lot of redundancy. Compression programs allow you to make files take less disk space by taking out that redundancy. If there is no redundancy, you won't gain much by using compression. In all examples of the **tar** command that you have seen so far, not a single byte has been compressed. Originally, after creating the archive, it had to be compressed with a separate compression utility, such as **gzip** or **bzip2**. After having created **home.tar**, you can compress it with **gzip** **home.tar**. **gzip** replaces **home.tar** with its compressed version, **home.tar.gz**, which takes significantly less space.

As an alternative to using **gzip**, you can use the **bzip2** utility. Originally, **bzip2** used a more efficient encryption algorithm, which resulted in smaller file sizes, but currently hardly any difference in file size exists between the result of **bzip2** and the result of **gzip**.

To decompress files that have been compressed with **gzip** or **bzip2**, you can use the **gunzip** and **bunzip2** utilities; you work with some examples of this command in Exercise 3-5.

As an alternative to using these utilities from the command line, you can include the **-z (gzip)** or **-j (bzip2)** option while creating the archive with **tar**. This will immediately compress the archive while it is created. There is no need to use these options while extracting. The **tar** utility will recognize the compressed content and automatically decompress it for you. In Exercise 3-5, you apply the newly acquired **tar** skills. Table 3-6 gives an overview of the most significant **tar** options.

**Key
Topic**

Table 3-6 Overview of **tar** Options

Option	Use
c	Creates an archive.
v	Shows verbose output while tar is working.
t	Shows the contents of an archive.
z	Compresses/decompresses the archive while creating it, by using gzip .
j	Compresses/decompresses the archive by using bzip2 .
x	Extracts an archive.
u	Updates an archive; only newer files will be written to the archive.
C	Changes the working directory before performing the command.
r	Appends files to an archive.

Exercise 3-5 Using tar

1. Open a root shell on your server. By logging in, the home directory of user root will become the current directory, so all relative filenames used in this exercise refer to /root/.
2. Type **tar -cvf etc.tar /etc** to archive the contents of the /etc directory.
3. Type **file etc.tar** and read the information that is provided by the command. This should look like the following:

```
[root@server1 ~]# file etc.taretc.
tar: POSIX tar archive (GNU)
```
4. Type **gzip etc.tar**.
5. Type **tar tvf etc.tar.gz**. Notice that the **tar** command has no issues reading from a **gzip** compressed file. Also notice that the archive content consists of all relative filenames.
6. Type **tar xvf etc.tar.gz etc/hosts**.
7. Type **ls -R**. Notice that a subdirectory etc has been created in the current directory. In this subdirectory, the file hosts has been restored.
8. Type **gunzip etc.tar.gz**. This decompresses the compressed file but does not change anything else with regard to the **tar** command.
9. Type **tar xvf etc.tar -C /tmp etc/passwd**. This extracts the password file to the /tmp directory.
10. Type **tar cjvf homes.tar /home**. This creates a compressed archive of the home directory to the home directory of user root.
11. Type **rm -f *gz *tar** to remove all files resulting from exercises in this chapter from the home directory of /root.

Summary

In this chapter, you learned how to work with essential file management tools. You learned how the Linux directory structure is organized by default, and you learned what file types to expect in which directories. You also learned how to find your way in the directory structure and to work with files.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 3-7 lists a reference of these key topics and the page number on which each is found.



Table 3-7 Key Topics for Chapter 3

Key Topic Element	Description	Page
Table 3-2	FHS overview	56
Table 3-3	Wildcard overview	62
Paragraph	Definition of an absolute filename	63
Paragraph	Definition of a relative filename	63
Table 3-4	ls common command-line options	64
Paragraph	Definition of an inode	68
Table 3-5	ln usage examples	69
Table 3-6	Overview of tar options	74

Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

File System Hierarchy Standard (FHS), directory, mount, device, folder, root directory, path, hard link, symbolic link, absolute filename, relative filename, inode, **tar**, **gzip**, compression, archiving, **bzip2**, star

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. Which directory would you go to if you were looking for configuration files?
2. Which command enables you to display a list of current directory contents, with the newest files listed first?
3. Which command enables you to rename the file `myfile` to `yourfile`?
4. Which command enables you to wipe an entire directory structure, including all of its contents?
5. How do you create a link to the directory `/tmp` in your home directory?
6. How would you copy all files that have a name that starts with `a`, `b`, or `c` from the directory `/etc` to your current directory?
7. Which command enables you to create a link to the directory `/etc` in your home directory?
8. What is the safe option to remove a symbolic link to a directory?
9. How do you create a compressed archive of the directories `/etc` and `/home` and write that archive to `/tmp/etchome.tgz`?
10. How would you extract the file `/etc/passwd` from `/tmp/etchome.tgz` that you have created in the previous step?

End-of-Chapter Lab

In this chapter, you learned how to perform basic file management tasks. Managing files is an essential task for a Linux administrator. This end-of-chapter lab enables you to practice these skills and make sure that you master them before taking the RHCSA exam.

Lab 3.1

1. Log in as user root. In the home directory of root, create one archive file that contains the contents of the /home directory and the /etc directory. Use the name /root/essentials.tar for the archive file.
2. Copy this archive to the /tmp directory. Also create a hard link to this file in the / directory.
3. Rename the file /essentials.tar to **/archive.tar**.
4. Create a symbolic link in the home directory of the user root that refers to /archive.tar. Use the name **link.tar** for the symbolic link.
5. Remove the file /archive.tar and see what happened to the symbolic link. Remove the symbolic link also.
6. Compress the /root/essentials.tar file.

This page intentionally left blank



The following topics are covered in this chapter:

- Using Common Text File–Related Tools
- A Primer to Using Regular Expressions
- Using **grep** to Analyze Text
- Working with Other Useful Text Processing Utilities

The following RHCSA exam objectives are covered in this chapter:

- Use **grep** and regular expressions to analyze text
- Create and edit text files

Since the early days of UNIX, working with text files has been an important administrator skill. Even on modern Linux versions such as Red Hat Enterprise Linux 8, working with text files is still an important skill. By applying the correct tools, you'll easily find the information you need. This chapter is about these tools. Make sure that you master them well, because good knowledge of these tools really will make your work as a Linux administrator a lot easier. At the same time, it will increase your chances of passing the RHCSA exam.

Working with Text Files

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 4-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 4-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Using Common Text File–Related Tools	1–5
A Primer to Using Regular Expressions	6–8
Using grep to Analyze Text	10
Working with Other Useful Text Processing Utilities	9

1. Which command was developed to show only the first ten lines in a text file?
 - a. **head**
 - b. **top**
 - c. **first**
 - d. **cat**
2. Which command enables you to count the number of words in a text file?
 - a. **count**
 - b. **list**
 - c. **ls -l**
 - d. **wc**

3. Which key on your keyboard do you use in **less** to go to the last line of the current text file?
- a. End
 - b. Page Down
 - c. q
 - d. G
4. Which option is missing (...) from the following command, assuming that you want to filter the first field out of the `/etc/passwd` file and assuming that the character that is used as the field delimiter is a `:`?
- ```
cut ... : -f 1 /etc/passwd
```
- a. -d
  - b. -c
  - c. -t
  - d. -x
5. Which option is missing (...) if you want to sort the third column of the output of the command **ps aux**?
- ```
ps aux | sort ...
```
- a. -k3
 - b. -s3
 - c. -k f 3
 - d. -f 3
6. Which of the following commands would only show lines in the file `/etc/passwd` that start with the text *anna*?
- a. `grep anna /etc/passwd`
 - b. `grep -v anna /etc/passwd`
 - c. `grep $anna /etc/passwd`
 - d. `grep ^anna /etc/passwd`
7. Which regular expression do you use to make the previous character optional?
- a. ?
 - b. .
 - c. *
 - d. &

8. Which regular expression is used as a wildcard to refer to any single character?
- a. `?`
 - b. `.`
 - c. `*`
 - d. `&`
9. Which command prints the fourth field of a line in the `/etc/passwd` file if the text *user* occurs in that line?
- a. `awk '/user/ { print $4 }' /etc/passwd`
 - b. `awk -d : '/user/ { print $4 }' /etc/passwd`
 - c. `awk -F : '/user/ $4' /etc/passwd`
 - d. `awk -F : '/user/ { print $4 }' /etc/passwd`
10. Which option would you use with **grep** to show only lines that do *not* contain the regular expression that was used?
- a. `-x`
 - b. `-v`
 - c. `-u`
 - d. `-q`

Foundation Topics

Using Common Text File–Related Tools

Before we start talking about the best possible way to find text files containing specific text, let's take a look at how you can display text files in an efficient way. Table 4-2 provides an overview of some common commands often used for this purpose.

**Key
Topic**

Table 4-2 Essential Tools for Managing Text File Contents

Command	Explanation
less	Opens the text file in a pager, which allows for easy reading
cat	Dumps the contents of the text file on the screen
head	Shows the first ten lines of the text file
tail	Shows the last ten lines of the text file
cut	Used to filter specific columns or characters from a text file
sort	Sorts the contents of a text file
wc	Counts the number of lines, words, and characters in a file

Apart from using these commands on a text file, they may also prove very useful when used in pipes. You can use the command **less /etc/passwd**, for example, to open the contents of the `/etc/passwd` file in the **less** pager, but you can also use the command **ps aux | less**, which sends the output of the command **ps aux** to the **less** pager to allow for easy reading.

Doing More with less

In many cases, as a Linux administrator you'll need to read the contents of text files. The **less** utility offers a convenient way to do so. To open the contents of a text file in **less**, just type **less** followed by the name of the file you want to see, as in **less /etc/passwd**.

From **less**, you can use the Page Up and Page Down keys on your keyboard to browse through the file contents. Seen enough? Then you can press **q** to quit **less**. Also very useful is that you can easily search for specific contents in **less** using **/sometext** for a forward search and **?sometext** for a backward search. Repeat the last search by using **n**.

If you think this sounds familiar, it should. You have seen similar behavior in **vim** and **man**. That is because all of these commands are based on the same code.

NOTE Once upon a time, **less** was developed because it offered more features than the classical UNIX tool **more** that was developed to go through file contents page by page. So, the idea was to do more with **less**. Developers did not like that, so they enhanced the features of the **more** command as well. The result is that both **more** and **less** offer many features that are similar, and which tool you use doesn't really matter that much anymore. There is one significant difference, though, and that is the **more** utility ends if the end of the file is reached. To prevent this behavior, you can start **more** with the **-p** option.

In Exercise 4-1, you apply some basic **less** skills to work with file contents and command output.

Exercise 4-1 Applying Basic **less** Skills

1. From a terminal, type **less /etc/passwd**. This opens the `/etc/passwd` file in the **less** pager.
2. Press **G** to go to the last line in the file.
3. Type **/root** to look for the text *root*. You'll see that all occurrences of the text *root* are highlighted.
4. Press **q** to quit **less**.
5. Type **ps aux | less**. This sends the output of the **ps aux** command (which shows a listing of all processes) to **less**. Browse through the list.
6. Press **q** to quit **less**.

Showing File Contents with **cat**

The **less** utility is useful to read long text files. If a text file is not that long, you are probably better off using **cat**. This tool just dumps the contents of the text file on the terminal it was started from. This is convenient if the text file is short. If the text file is long, however, you'll see all contents scrolling by on the screen, and only the lines that fit on the terminal screen are displayed. Using **cat** is simple. Just type **cat** followed by the name of the file you want to see. For instance, use **cat /etc/passwd** to show the contents of this file on your computer screen.

TIP The **cat** utility dumps the contents of a file to the screen from the beginning to the end, which means that for a long file you'll see the last lines of the file only. If you are interested in the first lines, you can use the **tac** utility, which gives the inversed result of **cat**.

Displaying the First or Last Lines of a File with **head** and **tail**

If a text file contains much information, it can be useful to filter the output a bit. You can use the **head** and **tail** utilities to do that. Using **head** on a text file will show by default the first ten lines of that file. Using **tail** on a text file shows the last ten lines by default. You can adjust the number of lines that are shown by adding **-n** followed by the number you want to see. So, **tail -n 5 /etc/passwd** shows the last five lines of the `/etc/passwd` file.

TIP With older versions of **head** and **tail**, you had to use the **-n** option to specify the number of lines you wanted to see. With current versions of both utilities, you may also omit the **-n** option. So, using either **tail -5 /etc/passwd** or **tail -n 5 /etc/passwd** gives you the exact same results.

Another useful option that you can use with **tail** is **-f**. This option starts by showing you the last ten lines of the file you've specified, but it refreshes the display as new lines are added to the file. This is convenient for monitoring log files. The command **tail -f /var/log/messages** (which has to be run as the root user) is a common command to show in real time messages that are written to the main log file `/var/log/messages`. To close this screen, press Ctrl-C.

By combining **tail** and **head**, you can do smart things as well. Suppose, for instance, that you want to see line number 11 of the `/etc/passwd` file. To do that, use **head -n 11 /etc/passwd | tail -n 1**. The command before the pipe shows the first 11 lines from the file. The result is sent to the pipe, and on that result **tail -n 1** is used, which leads to only line number 11 being displayed. In Exercise 4-2, you apply some basic **head** and **tail** operations to get the exact results that you want.

Exercise 4-2 Using Basic **head** and **tail** Operations

1. Type **tail -f /var/log/messages**. You'll see the last lines of `/var/log/messages` being displayed. The file doesn't close automatically.
2. Press Ctrl-C to quit the previous command.
3. Type **head -n 5 /etc/passwd** to show the first five lines in `/etc/passwd`.
4. Type **tail -n 2 /etc/passwd** to show the last two lines of `/etc/passwd`.
5. Type **head -n 5 /etc/passwd | tail -n 1** to show only line number 5 of the `/etc/passwd` file.

Filtering Specific Columns with cut

When working with text files, it can be useful to filter out specific fields. Imagine that you need to see a list of all users in the `/etc/passwd` file. In this file, several fields are defined, of which the first contains the name of the users who are defined.

To filter out a specific field, the **cut** command is useful. To do this, use the **-d** option to specify the field delimiter followed by **-f** with the number of the specific field you want to filter out. So, the complete command is **cut -d : -f 1 /etc/passwd** if you want to filter out the first field of the `/etc/passwd` file. You can see the result in Example 4-1.

Example 4-1 Filtering Specific Fields with cut

```
[root@localhost ~]# cut -f 1 -d : /etc/passwd
root
bin
daemon
adm
lp
sync
shutdown
halt
...
```

Sorting File Contents and Output with sort

Another very useful command to use on text files is **sort**. As you can probably guess, this command sorts text. If you type **sort /etc/passwd**, for instance, the content of the `/etc/passwd` file is sorted in byte order. You can use the **sort** command on the output of a command also, as in **cut -f 1 -d : /etc/passwd | sort**, which sorts the contents of the first column in the `/etc/passwd` file.

By default, the **sort** command sorts in byte order. Notice that this looks like alphabetical order, but it is not, as all capital letters are shown before lowercase letters. So Zoo would be listed before apple. In some cases, that is not convenient because the content that needs sorting may be numeric or in another format. The **sort** command offers different options to help sorting these specific types of data. Type, for instance, **cut -f 2 -d : /etc/passwd | sort -n** to sort the second field of the `/etc/passwd` file in numeric order. It can be useful also to sort in reverse order; if you use the command **du -h | sort -rn**, you get a list of files sorted with the biggest file in that directory listed first.

You can also use the **sort** command and specify which column you want to sort. To do this, use **sort -k3 -t : /etc/passwd**, for instance, which uses the field separator **:** to sort the third column of the **/etc/passwd** file.

You might also like the option to sort on a specific column of a file or the output of a command. An example is the command **ps aux**, which gives an overview of the busiest processes on a Linux server. (Example 4-2 shows partial output of this command.)

Example 4-2 Using **ps aux** to Find the Busiest Processes on a Linux Server

```
[root@localhost ~]# ps aux | tail -n 10
postfix 1350  0.0  0.7  91872  3848  ?      S   Jan24 0:00 qmgr -l
      -t unix -u
root      2162  0.0  0.3  115348 1928 tty1   Ss+ Jan24 0:00 -bash
postfix 5131  0.0  0.7  91804  3832  ?      S   12:10 0:00 pickup
      -l -t unix -u
root      5132  0.0  0.0  0          0      ?      S   12:10 0:00
      [kworker/0:1]
root      5146  0.0  0.9  133596 4868  ?      Ss  12:12 0:00 sshd:
      root@pts/0
root      5150  0.0  0.3  115352 1940 pts/0  Ss  12:12 0:00 -bash
root      5204  0.0  0.0  0          0      ?      S   12:20 0:00
      [kworker/0:2]
root      5211  0.0  0.0  0          0      ?      S   12:26 0:00
      [kworker/0:0]
root      5212  0.0  0.2  123356 1320 pts/0  R+  12:26 0:00 ps aux
root      5213  0.0  0.1  107928  672 pts/0  R+  12:26 0:00 tail -n 10
```

To sort the output of this command directly on the third column, use the command **ps aux | sort -k3**.

Counting Lines, Words, and Characters with wc

When working with text files, you sometimes get a large amount of output. Before deciding which approach to handling the large amount of output works best in a specific case, you might want to have an idea about the amount of text you are dealing with. In that case, the **wc** command is useful. In its output, this command gives three different results: the number of lines, the number of words, and the number of characters.

Consider, for example, the **ps aux** command. When executed as root, this command gives a list of all processes running on a server. One solution to count how many processes there are exactly is to pipe the output of **ps aux** through **wc**, as in **ps aux | wc**. You can see the result of the command in Example 4-3, which shows that the total number of lines is 90 and that there are 1,045 words and 7,583 characters in the command output.

Example 4-3 Counting the Number of Lines, Words, and Characters with **wc**

```
[root@localhost ~]# ps aux | wc
    90      1045      7583
```

A Primer to Using Regular Expressions

Working with text files is an important skill for a Linux administrator. You must know not only how to create and modify existing text files, but also how to find the text file that contains specific text.

It will be clear sometimes which specific text you are looking for. Other times, it might not. For example, are you looking for *color* or *colour*? Both spellings might give a match. This is just one example of why using flexible patterns while looking for text can prove useful. These flexible patterns are known as *regular expressions* in Linux.

To understand regular expressions a bit better, let's take a look at a text file example, shown in Example 4-4. This file contains the last six lines from the `/etc/passwd` file. (This file is used for storing Linux accounts; see Chapter 6, “User and Group Management,” for more details.)

Example 4-4 Example Lines from `/etc/passwd`

```
[root@localhost ~]# tail -n 6 /etc/passwd
anna:x:1000:1000::/home/anna:/bin/bash
rihanna:x:1001:1001::/home/rihanna:/bin/bash
annabel:x:1002:1002::/home/annabel:/bin/bash
anand:x:1003:1003::/home/anand:/bin/bash
joanna:x:1004:1004::/home/joanna:/bin/bash
joana:x:1005:1005::/home/joana:/bin/bash
```

Now suppose that you are looking for the user *anna*. In that case, you could use the general regular expression parser **grep** to look for that specific string in the file */etc/passwd* by using the command **grep anna /etc/passwd**. Example 4-5 shows the results of that command, and as you can see, way too many results are shown.

Example 4-5 Example of Why You Need to Know About Regular Expressions

```
[root@localhost ~]# grep anna /etc/passwd
anna:x:1000:1000::/home/anna:/bin/bash
rihanna:x:1001:1001::/home/rihanna:/bin/bash
annabel:x:1002:1002::/home/annabel:/bin/bash
joanna:x:1004:1004::/home/joanna:/bin/bash
```

**Key
Topic**

A regular expression is a search pattern that allows you to look for specific text in an advanced and flexible way.

Using Line Anchors

In Example 4-5, suppose that you wanted to specify that you are looking for lines that start with the text *anna*. The type of regular expression that specifies where in a line of output the result is expected is known as a *line anchor*.

To show only lines that start with the text you are looking for, you can use the regular expression **^** (in this case, to indicate that you are looking only for lines where *anna* is at the beginning of the line; see Example 4-6).

Example 4-6 Looking for Lines Starting with a Specific Pattern

```
[root@localhost ~]# grep ^anna /etc/passwd
anna:x:1000:1000::/home/anna:/bin/bash
annabel:x:1002:1002::/home/annabel:/bin/bash
```

Another regular expression that relates to the position of specific text in a specific line is **\$**, which states that the line ends with some text. For instance, the command **grep ash\$ /etc/passwd** shows all lines in the */etc/passwd* file that end with the text *ash*. This command shows all accounts that have a shell and are able to log in (see Chapter 6 for more details).

Using Escaping in Regular Expressions

Although not mandatory, when using regular expressions, it is a good idea to use escaping to prevent regular expressions from being interpreted by the shell. When a command line is entered, the Bash shell parses the command line, looking for any

special characters like `*`, `$`, and `?`. It will next interpret these characters. The point is that regular expressions use some of these characters as well, and to make sure the Bash shell doesn't interpret them, you should use escaping.

In many cases, it is not really necessary to use escaping; in some cases, the regular expression fails without escaping. To prevent this from ever happening, it is a good idea to put the regular expression between quotes. So, instead of typing `grep ^anna /etc/passwd`, it is better to use `grep '^anna' /etc/passwd`, even if in this case both examples work.

Using Wildcards and Multipliers

In some cases, you might know which text you are looking for, but you might not know how the specific text is written. Or you might just want to use one regular expression to match different patterns. In those cases, wildcards and multipliers come in handy.

To start with, there is the dot (`.`) regular expression. This is used as a wildcard character to look for one specific character. So, the regular expression `r.t` would match the strings `rat`, `rot`, and `rut`.

In some cases, you might want to be more specific about the characters you are looking for. If that is the case, you can specify a range of characters that you are looking for. For instance, the regular expression `r[ou]t` matches the strings `rat`, `rot`, and `rut`.

Another useful regular expression is the multiplier `*`. This matches zero or more of the previous character. That does not seem to be very useful, but indeed it is, as you will see in the examples at the end of this section.

If you know exactly how many of the previous character you are looking for, you can specify a number also, as in `re{2\}d`, which would match `reed`, but not `red`. The last regular expression that is useful to know about is `?`, which matches zero or one of the previous character. Table 4-3 provides an overview of the most important regular expressions.

Table 4-3 Most Significant Regular Expressions

Regular Expression	Use
<code>^text</code>	Matches line that starts with specified text.
<code>text\$</code>	Matches line that ends with specified text.
<code>.</code>	Wildcard. (Matches any single character.)
<code>[abc]</code>	Matches a, b, or c.
<code>*</code>	Matches zero to an infinite number of the previous character.
<code>\{2\}</code>	Matches exactly two of the previous character.

**Key
Topic**

Regular Expression	Use
\{1,3\}	Matches a minimum of one and a maximum of three of the previous character.
colou?r	Matches zero or one of the previous character. This makes the previous character optional, which in this example would match both <i>color</i> and <i>colour</i> .

Let's take a look at an example of a regular expression that comes from the man page `semanage-fcontext` and relates to managing SELinux (see Chapter 22, “Managing SELinux”). The example line contains the following regular expression:

```
"/web(/.*)"?
```

In this regular expression, the text `/web` is referred to. This text string can be followed by the regular expression `(/.)?`, which means zero or one `(/.)`, which in fact means that it can be followed by nothing or `(/.)`. The `(/.)` refers to a slash, which may be followed by an unlimited number of characters. To state it differently, the regular expression refers to the text `/web`, which may or may not be followed by any characters.

Using grep to Analyze Text

The ultimate utility to work with regular expressions is **grep**, which stands for “general regular expression parser.” Quite a few examples that you have seen already were based on the **grep** command. The **grep** command has a couple of useful options to make it even more efficient. Table 4-4 describes some of the most useful options.

Key Topic

Table 4-4 Most Useful **grep** Options

Option	Use
-i	Not case sensitive. Matches upper- and lowercase letters.
-v	Only shows lines that do <i>not</i> contain the regular expression.
-r	Searches files in the current directory and all subdirectories.
-e	Searches for lines matching more than one regular expression.
-A <number>	Shows <number> of lines after the matching regular expression.
-B <number>	Shows <number> of lines before the matching regular expression.

In Exercise 4-3, you work through some examples using these **grep** options.

Exercise 4-3 Using Common grep Options

1. Type **grep '^#' /etc/sysconfig/sshd**. This shows that the file `/etc/sysconfig/sshd` contains a number of lines that start with the comment sign, `#`.
2. To view the configuration lines that really matter, type **grep -v '^#' /etc/sysconfig/sshd**. This shows only lines that do not start with a `#`.
3. Type **grep -v '^#' /etc/sysconfig/sshd -B 5**. This shows lines that do not start with a `#` sign but also the five lines that are directly before each of those lines, which is useful because in the preceding lines you'll typically find comments on how to use the specific parameters. However, you'll also see that many blank lines are displayed.
4. Type **grep -v -e '^#' -e '^\$' /etc/sysconfig/sshd**. This excludes all blank lines and lines that start with `#`.

Working with Other Useful Text Processing Utilities

The **grep** utility is a powerful utility that allows you to work with regular expressions. It is not the only utility, though. Some even more powerful utilities exist, like **awk** and **sed**, both of which are extremely rich and merit a book by themselves. The utilities were developed in the time that computers did not commonly have screens attached, and for that reason they do a good job in treating text files in a scripted way.

As a Linux administrator in the twenty-first century, you do not have to be a specialist in using these utilities anymore. It does make sense, however, to know how to perform some common tasks using these utilities. The most useful use cases are summarized in the following examples.

This command shows the fourth line from `/etc/passwd`:

```
awk -F : '{ print $4 }' /etc/passwd
```

This is something that can be done by using the **cut** utility as well, but the **awk** utility is more successful in distinguishing the fields that are used in command output of files. The bottom line is that if **cut** does not work, you should try the **awk** utility.

You can also use the **awk** utility to do tasks that you might be used to using **grep** for. Consider the following example:

```
awk -F : '/user/ { print $4 }' /etc/passwd
```


This command searches the `/etc/passwd` file for the text *user* and will print the fourth field of any matching line.

In this example, the “stream editor” **sed** is used to print the fifth line from the `/etc/passwd` file:

```
sed -n 5p /etc/passwd
```

The **sed** utility is a very powerful utility for filtering text from text files (like **grep**), but it has the benefit that it also allows you to apply modifications to text files, as shown in the following example:

```
sed -i s/old-text/new-text/g ~/myfile
```

In this example, the **sed** utility is used to search for the text *old-text* in `~/myfile` and replace all occurrences with the text *new-text*. Notice that the default **sed** behavior is to write the output to `STDOUT`, but the option **-i** will write the result directly to the file. Make sure that you know what you are doing before using this command, because it might be difficult to revert file modifications that are applied in this way.

You’ll like the following example if you’ve ever had a utility containing a specific line in a file that was erroneous:

```
sed -i -e '2d' ~/myfile
```

With this command, you can delete a line based on a specific line number. You can also make more complicated references to line numbers. Use, for instance, **sed -i -e '2d;20,25d' ~/myfile** to delete lines 2 and 20 through 25 in the file `~/myfile`.

TIP Do not focus on **awk** and **sed** too much. These are amazing utilities, but many of the things that can be accomplished using them can be done using other tools as well. The **awk** and **sed** tools are very rich, and you can easily get lost in them if you are trying to dig too deep.

Summary

In this chapter, you learned how to work with text files. You acquired some important skills like searching text files with **grep** and displaying text files or part of them with different utilities. You have also learned how regular expressions can be used to make the search results more specific. Finally, you learned about the very sophisticated utilities **awk** and **sed**, which allow you to perform more advanced operations on text files.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 4-5 lists a reference of these key topics and the page number on which each is found.

Table 4-5 Key Topics for Chapter 4

Key Topic Element	Description	Page
Table 4-2	Essential tools for managing text file contents	84
Paragraph	Definition of regular expressions	90
Table 4-3	Most useful regular expressions	91
Table 4-4	Most useful grep options	92

**Key
Topic**

Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

regular expression, pager, escaping, wildcard, multiplier, line anchor

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. Which command enables you to see the results of the **ps aux** command in a way that you can easily browse up and down in the results?
2. Which command enables you to show the last five lines from `~/samplefile`?
3. Which command do you use if you want to know how many words are in `~/samplefile`?
4. After opening command output using **tail -f ~/mylogfile**, how do you stop showing output?
5. Which **grep** option do you use to exclude all lines that start with either a `#` or a `;`?
6. Which regular expression do you use to match one or more of the preceding character?
7. Which **grep** command enables you to see *text* as well as *TEXT* in a file?
8. Which **grep** command enables you to show all lines starting with *PATH*, as well as the five lines just before that line?
9. Which **sed** command do you use to show line 9 from `~/samplefile`?
10. Which command enables you to replace the word *user* with the word *users* in `~/samplefile`?

End-of-Chapter Lab

In this end-of-chapter lab, you work with some of the most significant text processing utilities.

Lab 4.1

1. Describe two ways to show line 5 from the `/etc/passwd` file.
2. How would you locate all text files on your server that contain the current IP address? Do you need a regular expression to do this?

3. You have just used the **sed** command that replaces all occurrences of the text *Administrator* with *root*. Your Windows administrators do not like that very much. How do you revert?
4. Assuming that in the **ps aux** command the fifth line contains information about memory utilization, how would you process the output of that command to show the process that has the heaviest memory utilization first in the results list?
5. Which command enables you to filter the sixth column of **ps aux** output?
6. How do you delete the sixth line from the file ~/myfile?



The following topics are covered in this chapter:

- Working on Local Consoles
- Using SSH and Related Utilities

The following RHCSA exam objectives are covered in this chapter:

- Access remote systems using SSH
- Log in and switch users in multiuser targets
- Boot, reboot, and shut down a system normally
- Securely transfer files between systems
- Configure key-based authentication for SSH

You have already learned how to log in on Linux from a graphical environment. In this chapter, you learn about some other methods to access a Linux shell and start working. You learn how to work from local consoles and from Secure Shell (SSH) to connect to Linux. You also learn how to perform some basic tasks from these environments.

Connecting to Red Hat Enterprise Linux 8

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 5-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 5-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Working on Local Consoles	1–6
Using SSH and Related Utilities	7–10

1. Which is the correct term for the description here?

“Used to refer to the physical screen you are currently looking at as a user”

- a. Terminal
- b. Console
- c. Shell
- d. Interface

2. Which is the correct term for the description here?

“The environment from which a shell is used where users can enter their commands”

- a. Terminal
- b. Console
- c. Shell
- d. Interface

3. Which is the correct term for the description here?
“The environment that offers a command line on which users type the commands they want to use”
 - a. Terminal
 - b. Console
 - c. Shell
 - d. Interface
4. Which device file is associated with the virtual console that is opened after using the Alt-F6 key sequence?
 - a. /dev/console6
 - b. /dev/tty6
 - c. /dev/vty6
 - d. /dev/pts/6
5. Which of the following methods will open a pseudo terminal device? (Choose two.)
 - a. Log in using an SSH session
 - b. Press Alt-F2 to open a new nongraphical login
 - c. Right-click the graphical desktop and select Open in Terminal
 - d. Enter your username and password on a nongraphical console
6. Sometimes a server reboot may be necessary to accomplish a task. Which of the following is *not* typically one of them?
 - a. To recover from serious problems such as server hangs and kernel panics
 - b. To apply kernel updates
 - c. To apply changes to kernel modules that are being used currently and therefore cannot be reloaded easily
 - d. To apply changes to the network configuration
7. Which of the following is true about remote access to Linux servers from a Windows environment?
 - a. Open a shell terminal on Windows and type **ssh** to access Linux servers remotely. The **ssh** command is available as a default part of the Windows operating system.
 - b. Configure Remote Access on Windows if you want to access Linux servers running the sshd process.

- c. Install the PuTTY program on Windows to access sshd services on Linux from Windows.
 - d. You cannot remotely access Linux machines from Windows.
- 8. What is the name of the file in which the public key fingerprint of the SSH servers you have connected to in the past are stored?
 - a. /etc/ssh/remote_hosts
 - b. /etc/ssh/known_hosts
 - c. ~/.ssh/remote_hosts
 - d. ~/.ssh/known_hosts
- 9. To allow graphical applications to be used through an SSH session, you can set a parameter in the /etc/ssh/ssh_config file. Using this parameter makes it unnecessary to use the **-X** command-line option each time an SSH session is initiated. Which of the following parameters should be used?
 - a. **Host ***
 - b. **TunnelX11 yes**
 - c. **ForwardX11 yes**
 - d. **Xclient yes**
- 10. Which of the following statements about key-based SSH authentication is true?
 - a. After creating the key pair, you need to copy the private key to the remote server.
 - b. Use **scp** to copy the public key to the remote server.
 - c. Use **ssh-copy-id** to copy the public key to the remote server.
 - d. Use **ssh-keygen** on the server to generate a key pair that matches the client keys.

Foundation Topics

Working on Local Consoles

You have already learned how to log in on Linux by using a graphical console. In this section, you learn some more about the possibilities you have while working from either a graphical Linux console or a text-based Linux console.

Key Topic

Before we get into details, it makes sense to highlight the difference between the words *console* and *terminal*. In this book, I follow the common notion of a console as the environment the user is looking at. That means that the console is basically what you see when you are looking at your computer screen.

A terminal is an environment that is opened on the console and provides access to a text shell, which is the command-line environment that can be used to type commands. A terminal can be offered through a window while using a graphical console, but it can also be opened as the only thing that you see in a textual console. This means that in a textual environment, the words *console* and *terminal* are more or less equivalent. In a graphical environment, they are not. Think of it like this: You can have multiple terminals open on a console, but you cannot have multiple consoles open in one terminal.

Logging In to a Local Console

Roughly, there are two ways to make yourself known to a Linux server. Sometimes you just sit at the Linux console and interactively log in from the login prompt that is presented. In other cases, a remote connection is established. The second part of this chapter is about logging in from a remote session; in this part, you learn how to work from a local console.

If a Linux server boots with a graphical environment (the so-called graphical target), you see a login prompt on which a username and a field in which to enter a password can be entered. Many Linux servers do not use a graphical environment at all, though, and just present a text-based console, as shown in Figure 5-1.

```
Red Hat Enterprise Linux 8.0 Beta (Ootpa)
Kernel 4.18.0-32.el8.x86_64 on an x86_64

Activate the web console with: systemctl enable --now cockpit.socket

server1 login: _
```

FIGURE 5-1 Logging In from a Text Console

To log in from a text console, you need to know which user account you should use. A user `root` is always available, but using this account to do your work is often not a good idea; the user `root` has no limitations to access the system and can therefore do a lot of damage. A small mistake can have a huge impact. Typically, it is a better idea to log in as one of the locally defined users, and there are many reasons to do so, including the following:

- It will make it more difficult to make critical errors.
- On many occasions, you will not need root permissions anyway.
- If you only allow access to normal users and not to root, it will force an attacker to guess three different things: the name of that specific user, the password of that user, and the root password.
- If you do need root access anyway, you can use the `su` - command from the local user environment to open a root shell. This opens one root shell only, while you will still be an ordinary user in all other parts of your current session.
- Use `sudo` to configure specific administration tasks for specific users only. See Chapter 6, “User and Group Management,” for more information.

Switching Between Terminals in a Graphical Environment

When working in a graphical environment, it is relatively easy to open several different working environments, such as different terminal windows in which you can work from a shell. In the upper-left part of the graphical interface, click Activities, and in the Search bar that appears, type **term**, which presents an icon to open a terminal. Because terminals are opened as a subshell, you do not have to log in to each terminal again, and will get access as the same user account that was originally used to log in to the graphical environment (see Figure 5-2).




FIGURE 5-2 Using Different Terminal Windows from the Graphical Environment

Working from a graphical environment is convenient. As an administrator, you can open several terminal windows, and in each terminal window you can use the **su -** command to open a shell in which you can work with a different user identity. This allows you to easily test features and see the results of these tests immediately. Exercise 5-1 guides you through a common scenario where you can do this and see how testing things from one terminal window while monitoring from another terminal window can be convenient.

Exercise 5-1 Working from Several Terminal Windows Simultaneously

1. Start your computer and make sure to log in as a non-root user account from the graphical login window that is presented. You should have a local user with the name *student* and the password *password* that you can use for this purpose.
2. Click **Activities**, and type **term**. Next, click the terminal icon to open a new terminal window.
3. Repeat step 2 so that you have two terminal windows that are open simultaneously.
4. From one of the terminal windows, type the command **su -** and enter the password of the root user. Then, type **tail -f /var/log/secure**. This opens a trace on the file */var/log/secure*, where you can monitor security events in real time.
5. From the other terminal windows, type **su -**. When asked for a password, you normally enter the password for the user root. Enter a wrong password.
6. Now look at the terminal where the trace on */var/log/secure* is still open. You will see that an error message has been written to this file.
7. Press Ctrl-C to close the **tail -f** session on the */var/log/secure* file.

Working with Multiple Terminals in a Nongraphical Environment

In the previous section, you learned how to work with multiple terminals in a graphical environment. This is relatively easy because you just have to open a new terminal window. In a nongraphical environment, you only have one terminal interface that is available and that makes working in different user shell environments a bit more difficult.

To offer an option that makes working from several consoles on the same server possible, Linux uses the concept of a *virtual terminal*. This feature allows you to open six different terminal windows from the same console at the same time and use key sequences to navigate between them. To open these terminal windows, you can

use the key sequences Alt-F1 through Alt-F6. The following virtual consoles are available:

- **F1:** Gives access to the GNOME Display Manager (GDM) graphical login
- **F2:** Provides access to the current graphical console
- **F3:** Gives access back to the current graphical session
- **F4–F6:** Gives access to nongraphical consoles

TIP A convenient alternative to using the Alt-Function key sequences is offered by the **chvt** command. This command enables you to switch to a different virtual environment directly from the current environment. If you are in a graphical console right now, open a terminal and type **chvt 4**. This brings you to a login prompt on virtual terminal 3. Switch back to the graphical environment using the **chvt 3** command, or use **chvt 1** to switch back to a graphical login prompt.

Of these virtual consoles, the first one is used as the default console. It is commonly known as the *virtual console tty1*, and it has a corresponding device file in the `/dev` directory that has the name `/dev/tty1`. The other virtual consoles also have corresponding device files, which are numbered `/dev/tty1` through `/dev/tty6`.

When working from a graphical environment, it is also possible to open different virtual consoles. Because the combinations between the Alt key and the Function keys typically already have a meaning in the graphical environment, you need to use Ctrl-Alt-Function key instead. So, do not use Alt-F4 to open `/dev/tty4` from a graphical environment, but instead use Ctrl-Alt-F4. To get back to the graphical console, you can use the Alt-F3 key sequence. The Alt-F6 and Ctrl-Alt-F6 key sequences are essentially the same. It is important to use the Ctrl key as well when going from a GUI to a text environment. To go back from the text environment to the GUI environment, using the Ctrl key is optional.

NOTE A long time ago, big central computers were used, to which dumb terminal devices were connected. These dumb terminal devices consisted of nothing more than a monitor and keyboard attached to it. From each of these dumb terminal devices, a console session to the operating system could be started. On a modern Linux server, no dumb terminals are attached. They have been replaced with the virtual terminals, as described in the previous section.

Understanding Pseudo Terminal Devices

Every terminal used in a Linux environment has a device file associated with it. You've just learned that terminals that are started in a nongraphical environment are typically referred to through the devices `/dev/tty1` through `/dev/tty6`.

For terminal windows that are started from a graphical environment, pseudo terminals are started. These pseudo terminals are referred to using numbers in the `/dev/pts` directory. So, the first terminal window that is started from a graphical environment appears as `/dev/pts/1`, the second terminal windows appears as `/dev/pts/2`, and so on. In Exercise 5-2, you learn how to work with these pseudo terminal devices and see which user is active on which pseudo terminal.

NOTE On earlier versions of Linux, pseudo terminals were seen as `pty` devices. These types of terminals are now deprecated and replaced with the `pts` terminal types, as described before.

Exercise 5-2 Working with Pseudo Terminals

1. Log in to the graphical console, using a non-root user account.
2. Open a terminal window.
3. From the terminal window, type `w`. This will give an overview of all users who are currently logged in. Notice the column that mentions the `tty` the users are on, in which you see `pts/0` that refers to the terminal window.
4. Open another graphical terminal window. Type `su -` to become root.
5. Type `w` to display once more an overview of all users who are currently logged in. Notice that the root terminal shows as owned by the non-root user as well.

At this point, you know how to work with the console, terminals, virtual terminals, and pseudo terminals. In the section “Using SSH and Related Utilities” later in this chapter, you use SSH to open terminal sessions to your server. These sessions show as pseudo terminals as well.

Booting, Rebooting, and Shutting Down Systems

As an administrator of a Linux server, you occasionally have to reboot the Linux server. Rebooting a server is not often a requirement, but it can make your work a lot easier because it will make sure that all processes and tasks that were running on your server have re-read their configurations and initialized properly.

TIP Rebooting a Linux server is an important task on the RHCSA exam. Everything you have configured should still be working after the server has rebooted. So, make sure that you reboot at least once during the exam, but also after making critical modifications to the server configuration. If your server cannot reboot anymore after applying critical modifications to your server's configuration, at least you know where to look to fix the issues.

For an administrator who really knows Linux very thoroughly, rebooting a server is seldom necessary. Experienced administrators can often trigger the right parameter to force a process to reread its configurations. There are some scenarios, though, in which even experienced Linux administrators have to reboot:

- To recover from serious problems such as server hangs and kernel panics
- To apply kernel updates
- To apply changes to kernel modules that are being used currently and therefore cannot be reloaded easily

**Key
Topic**

When a server is rebooted, all processes that are running need to shut down properly. If the server is just stopped by pulling the power plug, much data will typically be lost. That is because processes that have written data do not typically write that data directly to disk, but instead store it in memory buffers (cache) from where it is committed to disk when it is convenient for the operating system.

To issue a proper reboot, the `systemd` process has to be alerted. The `systemd` process is the first process that was started when the server was started, and it is responsible for managing all other processes, directly or indirectly. As a result, on system reboots or halts, the `systemd` process needs to make sure that all these processes are stopped. To tell the `systemd` process this has to happen, a few commands can be used:

- **`systemctl reboot` or `reboot`**
- **`systemctl halt` or `halt`**
- **`systemctl poweroff` or `poweroff`**

When stopping a machine, you can use the **`systemctl halt`** command or the **`systemctl poweroff`** command. The difference between these two commands is that the **`systemctl poweroff`** command talks to power management on the machine to shut off power on the machine. This often does not happen when using **`systemctl halt`**.

NOTE Using the methods that have just been described will normally reboot or stop your machine. In some cases, these commands might not work. If that is the case, there is an emergency reset option as well. Using this option may prove useful if the machine is not physically accessible. To force a machine to reset, as root you can type **echo b > /proc/sysrq-trigger**. This command immediately resets the machine without saving anything. Notice that this command should be used only if there are no other options!

Using SSH and Related Utilities

In the previous sections in this chapter, you learned how to access a terminal if you have direct access to the server console. Many administrators work with servers that are not physically accessible. To manage these servers, Secure Shell (SSH) is normally used. In this section, you learn how to work with SSH.

On modern Linux distributions, Secure Shell is the common method to gain access to other machines over the network. In SSH, cryptography is used to ensure that you can be sure that you are connecting to the intended server. Also, traffic is encrypted while transmitted.

Accessing Remote Systems Using SSH

To access a server using SSH, two components are needed. On the remote server that you want to access, the **sshd** service must be running and offering services at port 22, and it should not be blocked by the firewall. After installation, Red Hat Enterprise Linux starts the **sshd** process automatically, and by default it is not blocked by the firewall.

If the SSH port is open, you can access it using the **ssh** command from the command line. The **ssh** command by default tries to reach the **sshd** process on the server port 22. If you have configured the **sshd** process to offer its services on a different port, use **ssh -p** followed by the port number you want to connect to.

The **ssh** command is available on all Linux distributions, and on Apple Mac computers as well, where it can be launched from a Mac terminal.

If you have a Windows version that does not have the Windows subsystem for Linux, the **ssh** command is not a native part of the Windows operating system. If you want to access Linux servers through SSH from a Windows computer, you need to install an SSH client like PuTTY on Windows. From PuTTY, different types of remote sessions can be established with Linux machines. Alternative SSH clients for Windows are available as well, such as MobaXterm, KiTTY, mRemoteNG, Bitvise, and Xshell.

Accessing another Linux machine from a Linux terminal is relatively easy. Just type **ssh** followed by the name of the other Linux machine. After connecting, you will be prompted for a password if a default configuration is used. This is the password of a user account with the same name of your current user account, but who should exist on the remote machine.

When remotely connecting to a Linux server, the SSH client tries to do that as the user account you are currently logged in with on the local machine. If you want to connect using a different user account, you can specify the name of this user on the command line, in the `user@server` format. If, for instance, you want to establish an SSH session as user `root` to a remote server, type **ssh root@remoteserver**. An alternative way of specifying the user account is by using the option **-l username**. So, **ssh remoteserver -l root** will also try to establish a root session to the remote server. In Exercise 5-3, you learn how to log in to a remote server using SSH.

Exercise 5-3 Using SSH to Log In to a Remote Server

This exercise assumes that a remote server is available and reachable. In this exercise, `server1` is used as the local server, and `server2` is the remote server on which the **sshd** process should be up and running. If you cannot access a remote server to perform the steps in the exercise, you might alternatively replace `server2` with `localhost`. It is obvious that by doing so you will not log in to a remote server, but you still use the **ssh** command to connect to an **sshd** process, and you'll get the full experience of working with **ssh**.

1. Open a root shell on `server2`. Type **systemctl status sshd**. This should show you that the **sshd** process is currently up and running.
2. To avoid any firewall-related problems, type **systemctl stop firewalld**.
3. Type **ip a | grep 'inet '**. (Notice the space between **inet** and the closing quote mark.) Notice the IPv4 address your server is currently using. In the rest of this exercise, it is assumed that `server2` is using IP address `192.168.4.220`. Replace that address with the address that you have found here.
4. Open a shell as a nonprivileged user on `server1`.
5. On `server1`, type **ssh 192.168.4.220 -l root**. This connects to the **sshd** process on `server2` and opens a root shell.
6. Before being prompted for a password, you see a message indicating that the authenticity of host `192.168.4.220` cannot be established (see Example 5-1). This message is shown because the host you are connecting to is not yet known on your current host, which might involve a security risk. Type **yes** to continue.

7. When prompted, enter the root password. After entering it, you now are logged in to server2.
8. Type **w**. Notice that the SSH session you have just opened shows as just another pseudo terminal session, but you'll see the source IP address in the FROM column.
9. Type **exit** to close the SSH session.

Example 5-1 Security Message Displayed When Logging In to a Remote Server for the First Time

```
[root@server1 ~]# ssh 192.168.4.220 -l root
The authenticity of host '192.168.4.220 (<no hostip for proxy
command>)' can't be established.
ECDSA key fingerprint is 35:64:36:f8:ac:4f:8a:94:aa:6e:4b:85:ed:76:0a:eb.
Are you sure you want to continue connecting (yes/no)?
```

NOTE On some occasions, using **ssh** to get access to a server will be slow. If you want to know why, use the **-v** option with the **ssh** command. This will start SSH in verbose mode and show all the individual components that are contacted. By doing so, you might get an indication why your server is being slow.

The security message is displayed in Example 5-1 because the remote server has never been contacted before and therefore there is no way to verify the identity of the remote server. After connecting to the remote server, a public key fingerprint is stored in the file `~/.ssh/known_hosts`.

The next time you connect to the same server, this fingerprint is checked with the encryption key that was sent over by the remote server to initialize contact. If the fingerprint matches, you will not see this message anymore.

In some cases, the remote host key fingerprint does not match the key fingerprint that is stored locally. That is a potentially dangerous situation. Instead of being connected to the intended server, you might be connected to the server of an evildoer. It does, however, also happen if you are connecting to an IP address that you have been connected to before but that is now in use by a different server. In that case, you just have to remove the key fingerprint from the `~/.ssh/known_hosts` file on the client computer. You can easily do so, using **sed**. For instance, use **sed -i -e '25d' ~/.ssh/known_hosts** to remove line 25 from the `known_hosts` file (assuming that is the line containing the erroneous key).

Using Graphical Applications in an SSH Environment

From an SSH session, by default you cannot start graphical applications. That is because of security; a remote host cannot display screens on your computer without specific permission to do that. There are two requirements for starting graphical applications through an SSH connection:

- An X server must be running on the client computer. The X server is the software component that creates the graphical screens.
- The remote host must be allowed to display screens on the local computer.

The easiest way to allow the remote host to display graphical screens on your computer is by adding the **-Y** option to the **ssh** command. So, use **ssh -Y linda@server2** if you want to connect as linda to server2, and also be able to start graphical applications.

As you have noticed, the **ssh** command gives you a few options. Table 5-2 shows some of the most common options available.

Table 5-2 Common **ssh** Options

Option	Use
-v	Verbose; shows in detail what is happening while establishing the connection
-Y	Enables support for graphical applications
-p <PORT>	Used to connect to an SSH service that is not listening on the default port 22

**Key
Topic**

As an administrator, you can also create a systemwide configuration that allows you to use *X forwarding*, which is starting graphical applications through an SSH session. As root, open the configuration file **/etc/ssh/ssh_config** and make sure it includes the following line:

```
ForwardX11 yes
```

The next time you use the **ssh** command, X forwarding will be available by default.

Securely Transferring Files Between Systems

If a host is running the **sshd** service, that service can also be used to securely transfer files between systems. To do that, you can use the **scp** command if you want the file to be copied, or **rsync** if you want to synchronize the file. Also, the **sftp** command is a part of the SSH solution and enables users to use an FTP command-line syntax to transfer files using **sshd**.

Using scp to Securely Copy Files

The **scp** command is similar to the **cp** command, which is used to copy local files, but it also includes an option that enables it to work with remote hosts. You can use **scp** to copy files and subdirectories to and from remote hosts, and subdirectories as well. To copy, for instance, the `/etc/hosts` file to the `/tmp` directory on `server2` using your current user account, use the following command:

```
scp /etc/hosts server2:/tmp
```

If you want to connect to `server2` as user `root` to copy the `/etc/passwd` file to your home directory, you use the following command:

```
scp root@server2:/etc/passwd ~
```

You can also use **scp** to copy an entire subdirectory structure. To do so, use the **-r** command, as in the following command:

```
scp -r server2:/etc/ /tmp
```

Notice that the **scp** command can be configured to connect to a nondefault SSH port also. It is a bit confusing, but to do this with the **scp** command, you need the **-P** option followed by the port number you want to connect to. Notice that **ssh** uses **-p** (lowercase) to specify the port it needs to connect to; the **scp** command uses an uppercase **-P**.

Using sftp to Securely Transfer Files

The **sftp** command provides an alternative to securely transfer files. Whereas the **scp** command provides an interface that is very similar to the **cp** command, the **sftp** command provides an FTP-like interface. Because even modern FTP servers are still transferring passwords and other sensitive data without using encryption, **sftp** should be considered as an alternative.

When working with **sftp**, you open an FTP client session to the remote server, where the only requirement on the remote server is that it should be running the **sshd** process. From the FTP client session, you use typical FTP client options, like **put** to upload a file or **get** to download a file.

Notice that when working with **sftp**, the local directory is important, even if after opening the FTP session you only see the remote directory on the server. When downloading a file, using the **get** command, the file will be stored in the current local directory, and when uploading a file using **put**, the file will be searched for in the local directory. Exercise 5-4 gives you a guided tour through using the **sftp** command and the **rsync** command, discussed next.

Using rsync to Synchronize Files

The **rsync** command uses SSH to synchronize files between a remote directory and a local directory. The advantage of synchronizing files is that only differences need to be considered. So, for example, if you synchronize a 100-MiB file in which only a few blocks have changed since the previous sync, only the changed blocks will be synchronized. This approach is also known as a delta sync.

When using the **rsync** command, multiple options are available. Table 5-3 provides an overview.

Table 5-3 Common **rsync** Options

Option	Use
-r	Synchronizes the entire directory tree
-l	Also synchronizes symbolic links
-p	Preserves symbolic links
-n	Performs only a dry run, not actually synchronizing anything
-a	Uses archive mode, thus ensuring that entire subdirectory trees and all file properties will be synchronized
-A	Uses archive mode, and in addition synchronizes ACLs
-X	Synchronizes SELinux context as well

Exercise 5-4 Using SSH to Log In to a Remote Server

1. From a terminal, type **sftp student@server2**. This gives you access to an SFTP prompt that is opened on server2.
2. Type **ls**. You'll see files in the current working directory on the remote server.
3. Type **pwd**. This shows the current directory on the remote server.
4. Type **lpwd**. This shows your local current directory.
5. Type **lcd /tmp**. This changes the local current directory to /tmp.
6. Now type **get <filename>**, where **<filename>** is the name of any file in the remote current directory. This file will be downloaded to the local /tmp directory.
7. Type **exit** to close your SFTP session.
8. Type **rsync -a server2:/etc/ /tmp**. This synchronizes the contents of the remote directory /etc in the local directory /tmp.
9. Type **ls -d /tmp/etc** to verify that the remote directory has been synchronized.

Configuring Key-Based Authentication for SSH

If SSH is used on the Internet, it might not be a good idea to allow password logins. To make SSH a bit more secure, it will always first try whether login using public/private keys is possible. Only if that is not possible is a password login used. The only thing you need to do to enable key-based login is to create a key pair; everything else is organized by default already.

When using public/private key-based authentication, the user who wants to connect to a server generates a public/private key pair. The private key needs to be kept private and will never be distributed. The public key is stored in the home directory of the target user on the SSH server in the file `.ssh/authorized_keys`.

When authenticating using key pairs, the user generates a hash derived from the private key. This hash is sent to the server, and if on the server it proves to match the public key that is stored on the server, the user is authenticated.

Using Passphrases or Not?

When creating a public/private key pair, you are prompted for a passphrase. If you want maximal security, you should enter a passphrase. You are prompted for that passphrase each time that you are using the private key to authenticate to a remote host. That is very secure, but it is not very convenient. To create a configuration that allows for maximal convenience, you can just press the Enter key twice to confirm that you do not want to set a passphrase. This is a typical configuration that is used for authentication between servers in a trusted environment where no outside access is possible anyway.

To create a key pair, use the **ssh-keygen** command. The **ssh-copy-id** command is next used to copy the public key over to the target server. In Exercise 5-5, you create a public/private key pair to log in to the server2 host. (If no remote host is available, you can use localhost as an alternative to verify the procedure.)

Exercise 5-5 Connecting to a Remote Server with Public/Private Keys

1. On server1, open a root shell.
2. Type **ssh-keygen**. When asked whether you want to use a passphrase, press Enter to use the passphrase-less setup.
3. When asked for the filename in which to store the (private) key, accept the default filename `~/.ssh/id_rsa`.
4. When asked to enter a passphrase, press Enter twice.

5. The private key now is written to the `~/.ssh/id_rsa` file and the public key is written to the `~/.ssh/id_rsa.pub` file.
6. Use **ssh-copy-id server2** to copy to server2 the public key you have just created. You are then asked for the password on the remote server one last time.
7. After copying the public key, verify that it can actually be used for authentication. To do this, type **ssh server2**. You should now authenticate without having to enter the password for the remote user account.

After copying the public key to the remote host, it will be written to the `~/.ssh/authorized_keys` file on that host. Notice that if multiple users are using keys to log in with that specific account, the `authorized_keys` file may contain a lot of public keys. Make sure never to overwrite it because that will wipe all keys that are used by other users as well!

Using the screen Command

You do not have to know about the **screen** command for the RHCSA exam, but you should know about it for your job because it enables you to open multiple terminal sessions, even if you are not in a graphical session. It also allows you to share a session with other users, as well as attach to and detach from remote terminal sessions. Before you can use **screen**, you need to install it. To do this, use the **yum install -y screen** command. Then, type **screen** to open a screen session. From the screen session, you can start any command you like.

The **screen** command is particularly useful when used from an SSH session. You can start a task that takes a long time from a screen session, detach from it, and attach to it later. The command continues running, even if you are going home and shut down your computer. The next day, you can easily attach to the screen session again to complete the task. To do this, just follow a simple procedure:

1. Open an SSH session.
2. From the SSH session, type **screen** to open a screen session.
3. Start whichever task you want to start and keep it running.
4. Use the Ctrl-A, D key sequence to detach from the screen session and log out from the SSH session.
5. When you are ready to reconnect, start the SSH session again. It is essential that you are using the same user account that you used before.
6. Attach to the screen session again using **screen -r**. You can now conveniently finish the work that you have started from the screen session before.

While working with **screen**, you are working in a specific application. From within this application you can issue specific commands. To use these specific screen commands, start by typing Ctrl-A, ?. This shows a list of all commands that are available. Every screen command is started with the Ctrl-A key sequence. An important command to remember is Ctrl-A, /, which will close the screen session. Make sure to remove screen sessions that you do not need anymore; otherwise, they will stay active until the next time you reboot your server!

If while trying to attach to a screen session you get a message that multiple screen sessions are currently running, you need to be more specific. To find out which screen sessions currently are running, use the **screen -ls** command. In Example 5-2, you can see how the **screen** command is used to show screen sessions that are currently in a detached state, and how the **screen -r** command gives a list of all currently available screen sessions, from which you can select the screen session you want to connect to by adding the number of the screen session to the **screen -r** command.

Example 5-2 Selecting the Right **screen** Session

```
[sander@lab ~]$ screen
[detached from 30500.pts-0.lab]
[sander@lab ~]$ screen
[detached from 30532.pts-0.lab]
[sander@lab ~]$ screen -r
There are several suitable screens on:
      30532.pts-0.lab (Detached)
      30500.pts-0.lab (Detached)
Type "screen [-d] -r [pid.]tty.host" to resume one of them.
[sander@lab ~]$ screen -r 30500
```

You can also use **screen** to work together with other users in the same session. To do this, follow these steps:

1. Make sure that both users are using SSH to connect to the system on which you want to work together. Both users must use the same user account when using SSH to connect.
2. One user needs to start a screen session, using the **screen** command.
3. The second user can just connect to the screen session using **screen -x**.

Summary

In this chapter, you learned how to connect to Red Hat Enterprise Linux 8. You learned the difference between consoles, terminals, and shells, and you learned how to set up terminal sessions locally as well as remotely. You also learned how to use SSH to connect to a remote server and how to securely copy files between servers. At the end of this chapter you learned how you can use the **screen** utility to start jobs in a session you can detach from and then reattach to later.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 5-4 lists a reference of these key topics and the page number on which each is found.

Table 5-4 Key Topics for Chapter 5

Key Topic Element	Description	Page
Paragraph	Definition of the words <i>console</i> and <i>terminal</i>	102
List	Situations that typically require a server reboot	107
Table 5-2	Common ssh options	111

**Key
Topic**

Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

console, terminal, tty, login shell, subshell, reboot, systemd, key-based login, public key, private key

Review Questions

The questions that follow use an open-ended format that is meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. What is the console?
2. On a server that currently has an operational graphical interface, you are at a text-based login prompt. Which key sequence do you use to switch back to your current work on the graphical interface?
3. Which command(s) show(s) all users that currently have a terminal session open to a Linux server?
4. On a server where no GUI is operational, what would you expect to be the device name that is used by the first SSH session that is opened to that server?
5. Which command would you use to get detailed information on what SSH is doing while logging in?
6. How do you initiate an SSH session with support for graphical applications?
7. What is the name of the configuration file that needs to be edited to modify SSH client settings?
8. How do you copy the `/etc/hosts` file to the directory `/tmp` on `server2` using the username `lisa`?
9. What is the name of the file in which public keys are stored for remote users who want to log in to this machine using key-based authentication?
10. Which command enables you to generate an SSH public/private key pair?

End-of-Chapter Labs

The end-of-chapter labs help you practice what you learned throughout the chapter. The first lab is about connecting to RHEL 8 locally, and the second lab is about using SSH to log in to a remote server.

Lab 5.1

1. Log in to the local console on `server1`. Make sure that `server1` does *not* show a graphical interface anymore, but just a text-based login prompt.
2. Log in from that environment and activate `tty6`. From `tty6`, switch back on the graphical interface and use the correct key sequence to go to the graphical interface.

Lab 5.2

1. Set up SSH-based authentication. From server2, use SSH to connect to server1.
2. Make sure that graphical applications are supported through the SSH session. Also set up key-based authentication so that no password has to be entered while connecting to the remote server.



The following topics are covered in this chapter:

- Different User Types
- Creating and Managing User Accounts
- Creating and Managing Group Accounts

The following RHCSA exam objectives are covered in this chapter:

- Create, delete, and modify local user accounts
- Change passwords and adjust password aging for local user accounts
- Create, delete, and modify local groups and group memberships
- Configure superuser access

On a Linux system, a wide variety of processes are normally being used. These processes need access to specific resources on the Linux system. To determine how these resources can be accessed, a difference is made between processes that run in kernel mode and processes that run without full permissions to the operating system. In the latter case user accounts are needed, not only to grant the required permissions to processes, but also to make sure that people can do their work. This chapter explains how to set up user and group accounts.

User and Group Management

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 6-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 6-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Different User Types	1–2
Creating and Managing User Accounts	3–6
Creating and Managing Group Accounts	7–10

1. Which statement about privileged users (root) is true?
 - a. A privileged user is a user who has access to a Linux system.
 - b. A privileged user with no access permissions can do nothing at all.
 - c. Privileged users are not restricted in any way.
 - d. On every server, at least one privileged user must be manually created while installing the server.
2. On a default installation of an RHEL 8 server, which group does the user typically need to be a member of to be able to use **sudo** to run all administration commands?
 - a. admin
 - b. root
 - c. sys
 - d. wheel

3. There are different ways that users can run tasks with root permissions. Which of the following is not one of them?
 - a. **sudo**
 - b. **runas**
 - c. **su**
 - d. PolicyKit
4. Which of the following is used to store the hash of the user's encrypted password?
 - a. **/etc/passwd**
 - b. **/etc/shadow**
 - c. **/etc/users**
 - d. **/etc/secure**
5. Which configuration file should you change to set the default location for all new user home directories?
 - a. **/etc/login.defaults**
 - b. **/etc/login.defs**
 - c. **/etc/default/useradd**
 - d. **/etc/default/login.defs**
6. Which command enables you to get information about password properties such as password expiry?
 - a. **chage -l**
 - b. **usermod --show**
 - c. **passwd -l**
 - d. **chage --show**
7. Which of the following files is not processed when a user starts a login shell?
 - a. **/etc/profile**
 - b. **/etc/.profile**
 - c. **~/.bashrc**
 - d. **~/.bash_profile**

8. Which utility can be used to edit group membership directly through the appropriate configuration file?
 - a. **vigr**
 - b. **vipw**
 - c. **vipasswd**
 - d. **usermod**
9. Which command can be used to list all groups a user is member of?
 - a. **userlist**
 - b. **grouplist**
 - c. **id**
 - d. **groups**
10. What can you do to ensure that no users, except for the user root, can log in temporarily?
 - a. Set the default shell to `/usr/sbin/nologin`
 - b. Set the default shell to `/bin/false`
 - c. Create a file with the name `/etc/nologin`
 - d. Create a file with the name `/etc/nologin.txt`

Foundation Topics

Different User Types

In this chapter, you learn how to create and manage user accounts. Before diving into the details of user management, you learn how users are used in a Linux environment.

Key Topic

Users on Linux

On Linux, there are two ways to look at system security. There are privileged users, and there are unprivileged users. The default privileged user is root. This user account has full access to everything on a Linux server and is allowed to work in system space without restrictions. The root user account is meant to perform system administration tasks and should be used for that only. For all other tasks, an unprivileged user account should be used.

To get information about a user account, you can use the **id** command. When using this command from the command line, you can see details about the current user. You can also use it on other user accounts to get details about those accounts. Example 6-1 shows an example of the output of the command.

Example 6-1 Getting More Information About Users with **id**

```
[root@localhost ~]# id linda
uid=1001(linda) gid=1001(linda) groups=1001(linda)
```

Working as Root

On all Linux systems, by default there is the user root, also known as the superuser. This account is used for managing Linux and has no restrictions at all. Root, for instance, can create other user accounts on the system. For some tasks, root privileges are required. Some examples are installing software, managing users, and creating partitions on disk devices. Generally speaking, all tasks that involve direct access to devices need root permissions.

Because the root account is so useful for managing a Linux environment, some people make a habit of logging in as root directly. That is not recommended, especially not when you are logging in to a graphical environment. When you log in as root in a graphical environment, all tasks that are executed run as root as well, and that involves an unnecessary security risk. Therefore, you should instead use one of the following alternative methods. Table 6-2 provides an overview of these methods.

Table 6-2 Methods to Run Tasks with Elevated Permissions**Key
Topic**

Method	Description
su	Opens a subshell as a different user, with the advantage that commands are executed as root only in the subshell
sudo	Enables you to set up an environment where specific tasks are executed with administrative privileges
PolicyKit	Enables you to set up graphical utilities to run with administrative privileges

Using su

From a terminal window, you can use the **su** command to start a subshell in which you have another identity. To perform administrative tasks, for instance, you can log in with a normal user account and type **su** to open a root shell. This brings the benefit that root privileges are used only in the root shell.

If you type just the command **su**, the username root is implied. But **su** can be used to run tasks as another user as well. Type **su linda** to open a subshell as the user linda, for example. When using **su** as an ordinary user, you are prompted for a password, and after entering that, you acquire the credentials of the target user:

```
[linda@localhost ~]$ su
Password:
[root@localhost linda]#
```

The subshell that is started when using **su** is an environment where you are able to work as the target user account, but environment settings for that user account have not been set. If you need complete access to the entire environment of the target user account, you can use **su -** to start a login shell. If you start a login shell, all scripts that make up the user environment are processed, which makes you work in an environment that is exactly the same as when logging in as that user.

TIP Using **su -** is better than using **su**. When the **-** is used, a login shell is started; without the **-**, some variables may not be set correctly. So, you are better off using **su -** immediately.

sudo

If a non-root user needs to perform a specific system administration task, the user does not need root access; instead, the system administrator can configure **sudo** to give that user administrator permissions to perform the specific task. The user then carries out the task by starting the command with **sudo** (and entering their password when prompted). So, instead of using commands like **useradd** as the root user, you can use an ordinary user account and type **sudo useradd**. This is definitely more secure because you will only be able to act as if you have administrator permissions while running this specific command.

When creating a Linux user during the installation process, you can select to grant administrator permissions to that specific user. If you select to do so, the user will be able to use all administrator commands using **sudo**. It is also possible to set up **sudo** privileges after installation by making the user a member of the group wheel. To do that in a very easy way, use this simple two-step procedure:

1. Make the administrative user account a member of the group wheel by using **usermod -aG wheel user**.
2. Type **visudo** and make sure the line **%wheel ALL=(ALL) ALL** is included.

Apart from this method, which would give a user access to all administrative commands, you can use **visudo** to edit the sudoers configuration file and give user access to specific commands only. For example, if you would include the line **linda ALL=/usr/bin/useradd, /usr/bin/passwd** in this file, that would allow user linda to run only the commands useradd and passwd with administrative privileges.

PolicyKit

Most administration programs with a graphical user interface use PolicyKit to authenticate as the root user. If a normal user who is not a member of the group wheel accesses such an application, he will be prompted for authentication. If a user who is a member of the group wheel opens a PolicyKit application, he will have to enter his own password. For the RHCSA exam, you do not have to know the details of PolicyKit. If you are interested, you can take a look at the man pages of the **pkexec** and **polkit** commands for more details.

In Exercise 6-1, you practice switching user accounts.

Exercise 6-1 Switching User Accounts

1. Log in to your system as a nonprivileged user and open a terminal.
2. Type **whoami** to see which user account you are currently using. Type **id** as well, and notice that you get more detail about your current credentials when using **id**.
3. Type **su**. When prompted for a password, enter the root password. Type **id** again. You see that you are currently root.
4. Type **visudo** and make sure that the line `%wheel ALL=(ALL) ALL` is included.
5. Type **useradd -G wheel laura** to create a user laura who is a member of the group wheel.
6. Type **id laura** to verify that the user laura has been added to the group wheel.
7. Set the password for laura by typing **passwd laura**. Enter the password **password** twice.
8. Log out and log in as laura.
9. Type **sudo useradd lori**. Enter the password when asked. Notice that user lori is created.

Managing User Accounts

Now that you know how to perform tasks as an administrative or nonadministrative user, it is time to learn how to manage user accounts on Linux. In this section, you learn what is involved.

System Accounts and Normal Accounts

A typical Linux environment has two kinds of user accounts. There are normal user accounts for the people who need to work on a server and who need limited access to the resources on that server. These user accounts typically have a password that is used for authenticating the user to the system. There are also system accounts that are used by the services the server is offering. Both types of user accounts share common properties, which are kept in the files `/etc/passwd` and `/etc/shadow`. Example 6-2 shows a part of the contents of the `/etc/passwd` file.

Example 6-2 Partial Contents of the /etc/passwd User Configuration File

```
ntp:x:38:38::/etc/ntp:/sbin/nologin
chrony:x:994:993::/var/lib/chrony:/sbin/nologin
abrt:x:173:173::/etc/abrt:/sbin/nologin
pulse:x:171:171:PulseAudio System Daemon:/var/run/pulse:/sbin/nologin
gdm:x:42:42::/var/lib/gdm:/sbin/nologin
gnome-initial-setup:x:993:991::/run/gnome-initial-setup:/sbin/nologin
postfix:x:89:89::/var/spool/postfix:/sbin/nologin
sshd:x:74:74:Privilege-separated SSH:/var/empty/sshd:/sbin/nologin
tcpdump:x:72:72:::/sbin/nologin
user:x:1000:1000:user:/home/user:/bin/bash
```

NOTE On many Linux servers, there are hardly any user accounts that are used by people. Many Linux servers are installed to run a specific service, and if people interact with that service, they will authenticate within the service.

As you can see in Example 6-2, to define a user account, different fields are used in /etc/passwd. The fields are separated from each other by a colon. The following is a summary of these fields, followed by a short description of their purpose.

**Key
Topic**

- **Username:** This is a unique name for the user. Usernames are important to match a user to her password, which is stored separately in /etc/shadow (see next bullet). On Linux, there can be no spaces in the username, and in general it's a good idea to specify usernames in all lowercase letters.
- **Password:** In the old days, the second field of /etc/passwd was used to store the hashed password of the user. Because the /etc/passwd file is readable by all users, this poses a security threat, and for that reason on current Linux systems the hashed passwords are stored in /etc/shadow (discussed in the next section).
- **UID:** Each user has a unique user ID (UID). This is a numeric ID. It is the UID that really determines what a user can do. When permissions are set for a user, the UID (and not the username) is stored in the file metadata. UID 0 is reserved for root, the unrestricted user account. The lower UIDs (typically up to 999) are used for system accounts, and the higher UIDs (from 1000 on by default) are reserved for people who need to connect a directory to the server. The range of UIDs that are used to create regular user accounts is set in /etc/login.defs.
- **GID:** On Linux, each user is a member of at least one group. This group is referred to as the *primary group*, and this group plays a central role in permissions management, as discussed later in this chapter. Users can be a member of additional groups, which are administered in the file /etc/groups.

- **Comment field:** The Comment field, as you can guess, is used to add comments for user accounts. This field is optional, but it can be used to describe what a user account is created for. Some utilities, such as the obsolete **finger** utility, can be used to get information from this field. The field is also referred to as the GECOS field, which stands for General Electric Comprehensive Operating System and had a specific purpose for identifying jobs in the early 1970s when General Electric was still an important manufacturer of servers.
- **Directory:** This is the initial directory where the user is placed after logging in, also referred to as the *home directory*. If the user account is used by a person, this is where the person would store his personal files and programs. For a system user account, this is the environment where the service can store files it needs while operating.
- **Shell:** This is the program that is started after the user has successfully connected to a server. For most users this will be **/bin/bash**, the default Linux shell. For system user accounts, it will typically be a shell like **/sbin/nologin**. The **/sbin/nologin** command is a specific command that silently denies access to users (to ensure that if by accident an intruder logs in to the server, he cannot get any shell access). Optionally, you can create an **/etc/nologin.txt** file, in which case only root will be able to log in but other users will see the contents of this file when their logins are denied.

A part of the user properties is stored in **/etc/passwd**, which was just discussed. Another part of the configuration of user properties is stored in the **/etc/shadow** file. The settings in this file are used to set properties of the password. Only the user root and processes running as root have access to **/etc/shadow**. Example 6-3 shows an example of **/etc/shadow** contents.

Example 6-3 Sample Content from **/etc/shadow**

```
[root@localhost ~]# tail -n 10 /etc/shadow
ntp:!!:16420:::::::
chrony:!!:16420:::::::
abrt:!!:16420:::::::
pulse:!!:16420:::::::
gdm:!!:16420:::::::
gnome-initial-setup:!!:16420:::::::
postfix:!!:16420:::::::
sshd:!!:16420:::::::
tcpdump:!!:16420:::::::
user:$6$3VZbGx1djo6FfyZo$/Trg7Q.3foIsIFYxBm6UnHuxxBrxQxHDnDuZxgS.
We/MAUhn8HboBZzpaMD8gfm.fmlB/ML9LnuaT7CbVXx31:16420:0:99999:7:::
```

The following fields are included in `/etc/shadow`:

**Key
Topic**

- **Login name:** Notice that `/etc/shadow` does not contain any UIDs, but usernames only. This opens a possibility for multiple users using the same UID but different passwords (which, by the way, is not recommended).
- **Encrypted password:** This field contains all that is needed to store the password in a secure way. If the field is empty, no password is set and the user cannot log in. If the field starts with an exclamation mark, login for this account currently is disabled.
- **Days since Jan. 1, 1970, that the password was last changed:** Many things on Linux refer to this date, which on Linux is considered the beginning of time. It is also referred to as *epoch*.
- **Days before password may be changed:** This allows system administrators to use a more strict password policy, where it is not possible to change back to the original password immediately after a password has been changed. Typically this field is set to the value 0.
- **Days after which password must be changed:** This field contains the maximal validity period of passwords. Notice in the last line of Example 6-3 that it is set to 99,999 (about 274 years), which is the default.
- **Days before password is to expire that user is warned:** This field is used to warn a user when a forced password change is upcoming. Notice in the last line of Example 6-3 that it is set to 7 days, which is the default (even if the password validity is set to 99,999 days).
- **Days after password expires that account is disabled:** Use this field to enforce a password change. After password expiry, the user no longer can log in. After the account has reached the maximum validity period, the account is locked out. This field allows for a grace period in which the user can change their password, but only during the login process. This field is set in days and is unset by default.
- **Days since Jan. 1, 1970, that account is disabled:** An administrator can set this field to disable an account on a specific date. This is typically a better approach than removing an account, as all associated properties and files of the account will be kept, but the account no longer can be used to authenticate on your server. Note that this field does not have a default value.
- **A reserved field, which was once added “for future use”:** This field was reserved a long time ago; it will probably never be used.

Most of the password properties can be managed with the `passwd` or `chage` command, which are discussed later in this chapter.

Creating Users

There are many solutions for creating users on a Linux server. To start, you can edit the contents of the `/etc/passwd` and `/etc/shadow` configuration files directly in an editor, using the **vi** command (with the risk of making an error that could make logging in impossible to anyone). Another option is to use **useradd**, which is the utility that you should use for creating users. To remove users, you can use the **userdel** command. Use **userdel -r** to remove a user and the complete user environment.

Modifying the Configuration Files

Creating a user account by modifying the configuration files simply requires adding one line to `/etc/passwd` and another line to `/etc/shadow`, in which the user account and all of its properties are defined. This method of creating users is not recommended, though. If you make an error, you might mess up the consistency of the file and make logging in completely impossible to anyone. Also, you might encounter locking problems if one administrator is trying to modify the file contents directly while another administrator wants to write a modification with some tool.

If you insist on modifying the configuration files directly, you should use **vi**. This command opens an editor interface on your configuration files, and more important, it sets the appropriate locks on the configuration files to prevent corruption. It does *not* check syntax, however, so make sure that you know what you are doing, because making even one typo might still severely mess up your server. If you want to use this tool to modify the `/etc/shadow` file, use **vi -s**. To edit the contents of the `/etc/group` file where groups are defined, a similar command with the name **vigr** exists.

NOTE It is nice to know that **vi** and **vigr** exist, but it is better not to use these utilities or anything else that opens the user and group configuration files directly. Instead, use tools like **useradd** and **groupmod**.

Using useradd

The **useradd** utility is probably the most common tool on Linux for managing users. It allows you to add a user account from the command line by using many of its parameters. Use, for instance, the command **useradd -m -u 1201 -G sales,ops linda** to create a user `linda` who is a member of the secondary groups `sales` and `ops` with UID 1201 and add a home directory to the user account as well. (Secondary groups are explained in the section “Understanding Linux Groups,” later in the chapter.)

Home Directories

All normal users will have a home directory. For people, the home directory is the directory where personal files can be stored. For system accounts, the home directory often contains the working environment for the service account.

As an administrator, you normally will not change home directory–related settings for system accounts because they are created automatically from the RPM post-installation scripts when installing the related software packages. If you have people who need user accounts, you probably do want to manage home directory contents a bit.

When creating home directories (which happens by default while creating users), the content of the “skeleton” directory is copied to the user home directory. The skeleton directory is `/etc/skel`, and it contains files that are copied to the user home directory at the moment this directory is created. These files will also get the appropriate permissions to ensure that the new user can use and access them.

By default, the skeleton directory contains mostly configuration files that determine how the user environment is set up. If in your environment specific files need to be present in the home directories of all users, you take care of that by adding the files to the skeleton directory.

Managing User Properties

For changing user properties, the same rules apply as for creating user accounts. You can either work directly in the configuration files using **vi** or use command-line tools.

The ultimate command-line utility for modifying user properties is **usermod**. It can be used to set all properties of users as stored in `/etc/passwd` and `/etc/shadow`, plus some additional tasks, such as managing group membership. There is just one task it does not do well: setting passwords. Although **usermod** has an option **-p** that tells you to “use encrypted password for the new password,” it expects you to do the password encryption before adding the user account. That does not make it particularly useful. If as root you want to change the user password, you’d use the **passwd** command.

Configuration Files for User Management Defaults

When working with tools such as **useradd**, some default values are assumed. These default values are set in two configuration files: `/etc/login.defs` and `/etc/default/useradd`. Example 6-4 shows the contents of `/etc/default/useradd`.

Example 6-4 `useradd` Defaults in `/etc/default/useradd`

```
[root@localhost skel]# cat /etc/default/useradd
# useradd defaults file
GROUP=100
HOME=/home
INACTIVE=-1
EXPIRE=
SHELL=/bin/bash
SKEL=/etc/skel
CREATE_MAIL_SPOOL=yes
```

As shown in Example 6-4, the `/etc/default/useradd` file contains some default values that are applied when using `useradd`.

In the file `/etc/login.defs`, different login-related variables are set. This file is used by different commands, and it relates to setting up the appropriate environment for new users. Here is a list of some of the most significant properties that can be set from `/etc/login.defs`:

- **MOTD_FILE:** Defines the file that is used as the “message of the day” file. In this file, you can include messages to be displayed after the user has successfully logged in to the server.
- **ENV_PATH:** Defines the `$PATH` variable, a list of directories that should be searched for executable files after logging in.
- **PASS_MAX_DAYS, PASS_MIN_DAYS, and PASS_WARN_AGE:** Define the default password expiration properties when creating new users.
- **UID_MIN:** The first UID to use when creating new users.
- **CREATE_HOME:** Indicates whether or not to create a home directory for new users.

**Key
Topic****Managing Password Properties**

You learned about the password properties that can be set in `/etc/shadow`. You can use two commands to change these properties for users: **chage** and **passwd**. The commands are rather straightforward, as long as you know what the options are used for. For instance, the command **passwd -n 30 -w 3 -x 90 linda** sets the password for user `linda` to a minimal usage period of 30 days and an expiry after 90 days, where a warning is generated 3 days before expiry.

Many of the tasks that can be accomplished with **passwd** can be done with **chage** also. For instance, use **chage -E 2020-12-31 bob** to have the account for user bob expire on December 31, 2020. To see current password management settings, use **chage -l** (see Example 6-5). The **chage** command also has an interactive mode; if you type **chage anna**, for instance, the command will prompt for all the password properties you want to set interactively.

Example 6-5 Showing Password Expiry Information with **chage -l**

```
linux:~ # chage -l linda
Last password change           : Apr 11, 2019
Password expires               : never
Password inactive              : never
Account expires                : never
Minimum number of days between password change : 0
Maximum number of days between password change : 99999
Number of days of warning before password expir : 7
```

Creating a User Environment

When a user logs in, an environment is created. The environment consists of some variables that determine how the user is working. One such variable, for instance, is **\$PATH**, which defines a list of directories that should be searched when a user types a command.

To construct the user environment, a few files play a role:



- **/etc/profile**: Used for default settings for all users when starting a login shell
- **/etc/bashrc**: Used to define defaults for all users when starting a subshell
- **~/.profile**: Specific settings for one user applied when starting a login shell
- **~/.bashrc**: Specific settings for one user applied when starting a subshell

When logging in, the files are read in this order, and variables and other settings that are defined in these files are applied. If a variable or setting occurs in more than one file, the last one wins.

In Exercise 6-2, you apply common solutions to create user accounts.

Exercise 6-2 Creating User Accounts

1. Type **vim /etc/login.defs** to open the configuration file `/etc/login.defs` and change a few parameters before you start creating users. Look for the parameter **CREATE_HOME** and make sure it is set to “yes.”
2. Use **cd /etc/skel** to go to the `/etc/skel` directory. Type **mkdir Pictures** and **mkdir Documents** to add two default directories to all user home directories. Also change the contents of the file `.bashrc` to include the line **export EDITOR=/usr/bin/vim**, which sets the default editor for tools that need to modify text files.
3. Type **useradd linda** to create an account for user linda. Then, type **id linda** to verify that linda is a member of a group with the name linda and nothing else. Also verify that the directories Pictures and Documents have been created in user linda’s home directory.
4. Use **passwd linda** to set a password for the user you have just created. Use the password **password**.
5. Type **passwd -n 30 -w 3 -x 90 linda** to change the password properties. This has the password expire after 90 days (**-x 90**). Three days before expiry, the user will get a warning (**-w 3**), and the password has to be used for at least 30 days before (**-n 30**) it can be changed.
6. Create a few more users: lisa, lori, and bob, using **for i in lisa, lori, bob; do useradd \$i; done**.
7. Use **grep lori /etc/passwd /etc/shadow /etc/group**. This shows the user lori created in all three critical files and confirms they have been set up correctly.

Creating and Managing Group Accounts

Every Linux user has to be a member of at least one group. In this section, you learn how to manage settings for Linux group accounts.

Understanding Linux Groups

Linux users can be a member of two different kinds of groups. First, there is the primary group. Every user must be a member of the primary group, and there is only one primary group. When creating files, the primary group becomes group owner of these files. (File ownership is discussed in detail in Chapter 7, “Permissions Management.”) Users can also access all files their primary group has access to.

The user's primary group membership is defined in `/etc/passwd`; the group itself is stored in the `/etc/group` configuration file.

Besides the mandatory primary group, users can be a member of one or more secondary groups as well. A user can be a member of a *secondary group* in addition to the primary group. Secondary groups are important to get access to files. If the group a user is a member of has access to specific files, the user will get access to those files also. Working with secondary groups is important, in particular in environments where Linux is used as a file server to allow people working for different departments to share files with one another. You have also seen how secondary group membership can be used to enable user administrative privileges through **sudo**, by making the user a member of the group `wheel`.

Creating Groups

As is the case for creating users, there are also different options for creating groups. The group configuration files can be modified directly using **vigr** or the command-line utility **groupadd**.

Creating Groups with vigr

With the **vigr** command, you open an editor interface directly on the `/etc/group` configuration file. In this file, groups are defined in four fields per group (see Example 6-6).

Example 6-6 Sample `/etc/group` Content

```
kvm:x:36:qemu
qemu:x:107:
libstoragemgmt:x:994:
rpc:x:32:
rpcuser:x:29:
"/etc/group.edit" 65L, 870C
```

The following fields are used in `/etc/group`:

- **Group name:** As is suggested by the name of the field, this contains the name of the group.
- **Group password:** Where applicable, this field contains a group password, a feature that is hardly used anymore. A group password can be used by users who want to join the group on a temporary basis, so that access to files the

group has access to is allowed. If a group password is used, it is stored in the `/etc/gshadow` file, as that file is root accessible only.

- **Group ID:** This field contains a unique numeric group identification number.
- **Members:** Here you find the names of users who are a member of this group as a secondary group. Note that this field does not show users who are a member of this group as their primary group.

As mentioned, in addition to `/etc/group`, there is the `/etc/gshadow` file. This file is not commonly used to store group passwords (because hardly anyone still uses them), but it does have a cool feature. In the third field of this file you can list administrators. This is a comma-separated list of users that can change passwords for group members, which are listed in the fourth field of this file. Note that specifying group members here is optional, but if it is done, the group member names must be the same as the group members in `/etc/group`.

Using groupadd to Create Groups

Another method to create new groups is by using the **groupadd** command. This command is easy to use. Just use **groupadd** followed by the name of the group you want to add. There are some advanced options; the only significant one is **-g**, which allows you to specify a group ID when creating the group.

Managing Group Properties

To manage group properties, **groupmod** is available. You can use this command to change the name or group ID of the group, but it does not allow you to add group members. To do this, you use **usermod**. As discussed before, **usermod -aG** will add users to new groups that will be used as their secondary group. Because a group does not have many properties, it is quite common that group properties are managed directly in the `/etc/group` file by using the **vigr** command.

In Exercise 6-3, you create two groups and add some users as members to these groups.

TIP Because a user's group membership is defined in two different locations, it can be difficult to find out which groups exactly a user is a member of. A convenient command to check this is **groupmems**. Use, for example, the command **groupmems -g sales -l** to see which users are a member of the group `sales`. This shows users who are a member of this group as a secondary group assignment, but also users who are a member of this group as the primary group assignment.

Exercise 6-3 Working with Groups

1. Type **groupadd sales** followed by **groupadd account** to add groups with the names sales and account.
2. Use **usermod** to add users linda and laura to the group sales, and lori and bob to the sales group account:

```
usermod -aG sales linda
usermod -aG sales laura
usermod -aG account lori
usermod -aG account bob
linux:~ # id linda
```

3. Type **id linda** to verify that user linda has correctly been added to the group sales. In the results of this command, you see that linda is assigned to a group with the name linda. This is her primary group and is indicated with the gid option. The **groups** parameter shows all groups she currently is a member of, which includes the primary group as well as the secondary group sales that she has just been assigned to.

```
uid=1000(linda) gid=1000(linda) groups=1000(linda),1001(sales)
```

Summary

In this chapter, you learned how to create users and groups. You learned which configuration files are used to store users and groups, and you learned which properties are used in these files. You also learned which utilities are available to manage user and group accounts.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have a couple of choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 6-3 lists a reference of these key topics and the page number on which each is found.

Table 6-3 Key Topics for Chapter 6

Key Topic Element	Description	Page
Section	Users on Linux	124
Table 6-2	Methods to run tasks with elevated permissions	125
List	Description of user account fields in <code>/etc/passwd</code>	128
List	Description of password property fields in <code>/etc/shadow</code>	130
List	Significant properties that can be set from <code>/etc/login.defs</code>	133
List	Files that play a role in constructing the user environment	134

**Key
Topic**

Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

user, password, GECOS, group, primary group, secondary group, privileged user, unprivileged user, root

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. What is the UID of user root?
2. What is the configuration file in which **sudo** is defined?
3. Which command should you use to modify a **sudo** configuration?
4. Which two files can be used to define settings that will be used when creating users?
5. How many groups can you create in `/etc/passwd`?
6. If you want to grant a user access to all admin commands through **sudo**, which group should you make that user a member of?
7. Which command should you use to modify the `/etc/group` file manually?
8. Which two commands can you use to change user password information?
9. What is the name of the file where user passwords are stored?
10. What is the name of the file where group accounts are stored?

End-of-Chapter Lab

You have now learned how to set up an environment where user accounts can log in to your server and access resources on your server. In this end-of-chapter lab, you learn how to configure an environment for users and groups.

Lab 6.1

Set up a shared group environment that meets the following requirements:

- Create two groups: sales and account.
- Create users bob, betty, bill, and beatrix. Make sure they have their primary group set to a private group that has the name of the user.
- Make bob and betty members of the group sales, and bill and beatrix members of the group account.
- Set a password policy that requires users to change their password every 90 days.

This page intentionally left blank



The following topics are covered in this chapter:

- Managing File Ownership
- Managing Basic Permissions
- Managing Advanced Permissions
- Managing ACLs
- Setting Default Permissions with **umask**
- Working with User-Extended Attributes

The following RHCSA exam objectives are covered in this chapter:

- List, set, and change standard ugo/rwx permissions
- Create and configure set-GID directories for collaboration
- Create and manage access control lists
- Diagnose and correct file permission problems

To get access to files on Linux, permissions are used. These permissions are assigned to three entities: the file owner, the group owner, and the others entity (which is everybody else). In this chapter, you learn how to apply permissions. The chapter starts with an overview of the basic partitions, after which the special permissions and access control lists (ACLs) are discussed. At the end of this chapter, setting default permissions through the **umask** is covered, as is managing user-extended attributes.

Permissions Management

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 7-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 7-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Managing File Ownership	1–3
Managing Basic Permissions	4
Managing Advanced Permissions	5
Managing ACLs	6–8
Setting Default Permissions with umask	9
Working with User-Extended Attributes	10

1. A user needs to work in a session where all new files that he creates will be group-owned by the group sales, until the session is closed. Which command would do that?
 - a. **chgrp sales**
 - b. **setgid sales**
 - c. **newgrp sales**
 - d. **setgroup sales**
2. Which command enables you to find all files on a system that are owned by user linda?
 - a. **find / -user linda**
 - b. **find / -uid linda**
 - c. **ls -l | grep linda**
 - d. **ls -R | find linda**

3. Which command does not set group ownership to the group sales for the file myfile?
 - a. **chgrp sales myfile**
 - b. **chown .sales myfile**
 - c. **chgrp myfile sales**
 - d. **chown :sales myfile**
4. Which command would be used to allow read and write permissions to the user and group owners and no permissions at all to anyone else?
 - a. **chown 007 filename**
 - b. **chmod 077 filename**
 - c. **chmod 660 filename**
 - d. **chmod 770 filename**
5. Which command enables you to set the SGID permission on a directory?
 - a. **chmod u+s /dir**
 - b. **chmod g-s /dir**
 - c. **chmod g+s /dir**
 - d. **chmod 1770 /dir**
6. You are trying to use the **setfacl** command to set ACLs on the directory /data, but you are getting an “operation not supported” message. Which of the following is the most likely explanation?
 - a. The **setfacl** command is not installed on your computer.
 - b. You are making an error typing the command.
 - c. The user or group to which you want to grant ACLs does not exist.
 - d. The file system lacks ACL support.
7. Which command grants rw permissions for the group sales to all new files that will be created in the /data directory and all of its subdirectories?
 - a. **setfacl -m d:g:sales:rw /data**
 - b. **setfacl -m d:g:sales:rwx /data**
 - c. **setfacl -R -m g:sales:rwx /data**
 - d. **setfacl -R -m g:sales:rw /data**

8. Which command enables you to make sure that others have no access to any new files that will be created in the /data directory, assuming that you want others to have read permissions on all other files?
- a. **setfacl -m d:o::- /data**
 - b. **setfacl -m o::- /data**
 - c. **umask 027 /data**
 - d. **umask 027**
9. Which of the following umask settings meets the following requirements:
- Grants all permissions to the owner of the file.
 - Grants read permissions to the group owner of the file.
 - Grants no permissions to others.
- a. 740
 - b. 750
 - c. 027
 - d. 047
10. Which command enables you to check all attributes that are currently set on myfile.
- a. **ls --attr myfile**
 - b. **getattr myfile**
 - c. **lsattr myfile**
 - d. **listattr myfile**

Foundation Topics

Managing File Ownership

Before we discuss permissions, you need to understand the role of file and directory ownership. File and directory ownership are vital for working with permissions. In this section, you first learn how you can see ownership. Then you learn how to change user and group ownership for files and directories.

Displaying Ownership

On Linux, every file and every directory has two owners: a user owner and a group owner. Apart from that, there is the “others” entity, which also is considered to be an entity to determine the permissions a user has. Collectively, the user, group, and others (ugo) owners are shown when listing permissions with the **ls -l** command.

These owners are set when a file or directory is created. On creation, the user who creates the file becomes the user owner, and the primary group of that user becomes the group owner. To determine whether you as a user have permissions to a file or a directory, the shell checks ownership. This happens in the following order:

1. The shell checks whether you are the user owner of the file you want to access, which is also referred to as the user of the file. If you are the user, you get the permissions that are set for the user, and the shell looks no further.
2. The shell checks whether you have obtained any permissions through user-assigned ACLs. If this is the case, these permissions are assigned and the shell looks no further.
3. If you are not the user owner and did not get permissions through ACLs, the shell checks whether you are a member of the group owner, which is also referred to as the group of the file. If you are a member of the group, you get access to the file with the permissions of the group, and the shell looks no further.
4. If you’ve come this far, the shell checks whether you have obtained any permissions through a group ACL. If that is the case, these permissions are applied, and the shell looks no further. If you are a member of multiple groups that have obtained permissions through ACLs, all of these permissions will apply.
5. If you are neither the user owner nor the group owner and have not obtained permissions through ACLs, you get the permissions of others entity.

To see current ownership assignments, you can use the **ls -l** command. This command shows the user owner and the group owner. In Example 7-1, you can see the ownership settings for directories in the directory /home.

Example 7-1 Displaying Current File Ownership

```
[root@server1 home]# ls -l
total 8
drwx-----. 3 bob bob          74 Feb 6 10:13 bob
drwx-----. 3 caroline caroline 74 Feb 6 10:13 caroline
drwx-----. 3 fozia fozia      74 Feb 6 10:13 fozia
drwx-----. 3 lara lara        74 Feb 6 10:13 lara
drwx-----. 5 lisa lisa       4096 Feb 6 10:12 lisa
drwx-----. 14 user user      4096 Feb 5 10:35 user
```

Using the **ls** command, you can display ownership for files in a given directory. It may on occasion be useful to get a list of all files on the system that have a given user or group as owner. To do this, you may use **find**. The **find** argument **-user** can be used for this purpose. For instance, the following command shows all files that have user linda as their owner:

```
find / -user linda
```

You can also use **find** to search for files that have a specific group as their owner. For instance, the following command searches all files that are owned by the group users:

```
find / -group users
```

Changing User Ownership

To apply appropriate permissions, the first thing to consider is ownership. To do this, you can use the **chown** command. The syntax of this command is not hard to understand:

```
chown who what
```

For instance, the following command changes ownership for the file account to user linda:

```
chown linda account
```

The **chown** command has a few options, of which one is particularly useful: **-R**. You might guess what it does, because this option is available for many other commands as well. It allows you to set ownership recursively, which allows you to set ownership of the current directory and everything below. The following command changes ownership for the directory /home and everything beneath it to user linda:

```
chown -R linda /home/linda
```

Changing Group Ownership

There are actually two ways to change group ownership. You can do it using **chown**, but there is also a specific command with the name **chgrp** that does the job. If you want to use the **chown** command, use a **.** or **:** in front of the group name. The following changes the group owner of directory `/home/account` to the group `account`:

```
chown .account /home/account
```

You can use **chown** to change user and/or group ownership in a number of ways, an overview of which follows:

- **chown lisa myfile** Sets user `lisa` as the owner of `myfile`
- **chown lisa.sales myfile** Sets user `lisa` as user owner and group `sales` as group owner of `myfile`
- **chown lisa:sales myfile** Sets user `lisa` as user owner and group `sales` as group owner of `myfile`
- **chown .sales myfile** Sets group `sales` as group owner of `myfile` without changing the user owner
- **chown :sales myfile** Sets group `sales` as group owner of `myfile` without changing the user owner

You can also use the **chgrp** command to change group ownership. Consider the following example, where you can use **chgrp** to set group ownership for the directory `/home/account` to the group `account`:

```
chgrp account /home/account
```

As is the case for **chown**, you can use the option **-R** with **chgrp** as well to change group ownership recursively.

Understanding Default Ownership

You might have noticed that when a user creates a file, default ownership is applied. The user who creates the file automatically becomes user owner, and the primary group of that user automatically becomes group owner. Normally, this is the group that is set in the `/etc/passwd` file as the user's primary group. If the user is a member of more groups, however, he can change the effective primary group so that new files will get the new primary group as group owner.

To show the current effective primary group, a user can use the **groups** command:

```
[root@server1 ~]# groups lisa
lisa : lisa account sales
```

If the current user *linda* wants to change the effective primary group, she can use the **newgrp** command, followed by the name of the group she wants to set as the new effective primary group. This group will be continued to be used as effective primary group until the user uses the **exit** command or logs out. Example 7-2 shows how user *linda* uses this command to make *sales* her effective primary group.

Example 7-2 Using **newgrp** to Change the Effective Primary Group

```
[lisa@server1 ~]$ groups
lisa account sales
[lisa@server1 ~]$ newgrp sales
[lisa@server1 ~]$ groups
sales lisa account
[lisa@server1 ~]$ touch file1
[lisa@server1 ~]$ ls -l
total 0
-rw-r--r--. 1 lisa sales 0 Feb 6 10:06 file1
```

After you change the effective primary group, all new files that the user creates will get this group as their group owner. To return to the original primary group setting, use **exit**.

To be able to use the **newgrp** command, a user has to be a member of that group. Alternatively, a group password can be set for the group using the **gpasswd** command, but that is uncommonly used. If a user uses the **newgrp** command but is not a member of the target group, the shell prompts for the group password. After you enter the correct group password, the new effective primary group is set.

Managing Basic Permissions

The Linux permissions system was invented in the 1970s. Because computing needs were limited in those years, the basic permission system that was created was rather limited as well. This basic permission system uses three permissions that can be applied to files and directories. In this section, you learn how the system works and how to modify these permissions.

Understanding Read, Write, and Execute Permissions

The three basic permissions allow users to read, write, and execute files. The effect of these permissions differs when applied to files or directories. If applied to a file, the read permission gives the right to open the file for reading. Therefore, you can read its contents, but it also means that your computer can open the file to do something with it. A program file that needs access to a library needs, for example,

read access to that library. From that follows that the read permission is the most basic permission you need to work with files.

If applied to a directory, read allows you to list the contents of that directory. You should be aware that this permission does not allow you to read files in the directory as well. The Linux permission system does not know inheritance, and the only way to read a file is by using the read permissions on that file. To open a file for reading, however, it is convenient to have read permissions to the directory because you would not see the file otherwise. Notice that without read permission on the directory, you can still open the file if you know that a file with that name exists. This is not very convenient and for that reason not commonly done.

As you can probably guess, the write permission, if applied to a file, allows you to write in the file. Stated otherwise, write allows you to modify the contents of existing files. It does not, however, allow you to create or delete new files. To do that, you need write permission on the directory where you want to create the file. To modify the permissions on a file, you don't need permissions on the file; you just have to be owner, or root. On directories, this permission also allows you to create and remove new subdirectories.

The execute permission is what you need to execute a file. Also, you need the execute permission on a directory if you want to do anything in that directory. It will never be set by default, which makes Linux almost completely immune to viruses. Only someone with ownership can apply the execute permission.

Whereas having the execute permission on files means that you are allowed to run a program file, if applied to a directory it means that you are allowed to use the `cd` command to go to that directory. This means that execute is an important permission for directories, and you will see that it is normally applied as the default permission to directories. Without it, there is no way to change to that directory! Table 7-2 summarizes the use of the basic permissions.

Key
Topic

Table 7-2 Use of Read, Write, and Execute Permissions

Permission	Applied to Files	Applied to Directories
Read	Open a file	List contents of directory
Write	Change contents of a file	Create and delete files
Execute	Run a program file	Change to the directory

Applying Read, Write, and Execute Permissions

To apply permissions, you use the `chmod` command. When using `chmod`, you can set permissions for user, group, and others. You can use this command in two modes: the relative mode and the absolute mode. In absolute mode, three digits are

used to set the basic permissions. The three digits apply to user, group, and others, respectively. Table 7-3 provides an overview of the permissions and their numeric representation.

Table 7-3 Numeric Representation of Permissions

Permission	Numeric Representation
Read	4
Write	2
Execute	1

**Key
Topic**

When setting permissions, calculate the value that you need. If you want to set read, write, and execute for the user, read and execute for the group, and read and execute for others on the file `/somefile`, for example, you use the following **chmod** command:

```
chmod 755 /somefile
```

When you use **chmod** in this way, all current permissions are replaced by the permissions you set. If you want to modify permissions relative to the current permissions, you can use **chmod** in relative mode. When using **chmod** in relative mode, you work with three indicators to specify what you want to do:

- First, you specify for whom you want to change permissions. To do this, you can choose between user (**u**), group (**g**), and others (**o**).
- Then, you use an operator to add or remove permissions from the current mode, or set them in an absolute way.
- At the end, you use **r**, **w**, and **x** to specify what permissions you want to set.

When changing permissions in relative mode, you may omit the “to whom” part to add or remove a permission for all entities. For instance, the following adds the execute permission for all users:

```
chmod +x somefile
```

When working in relative mode, you may use more complex commands as well. For instance, the following adds the write permission to the group and removes read for others:

```
chmod g+w,o-r somefile
```

In Exercise 7-1, you learn how to work with basic permissions by creating a directory structure for the groups that you have created earlier. You also assign the correct permissions to these directories.

TIP When using **chmod -R o+rx /data**, you set the execute permission on all directories as well as files in the /data directory. To set the execute permission to directories only, and not to files, use **chmod -R o+rX /data**. The uppercase X ensures that files will not get the execute permission unless the file has already set the execute permission for some of the entities. That makes X the more intelligent way of dealing with execute permissions; it will avoid setting that permission on files where it is not needed.

Exercise 7-1 Managing Basic Permissions

1. From a root shell, type **mkdir -p /data/sales /data/account**.
2. Before setting the permissions, change the owners of these directories using **chown linda.sales /data/sales** and **chown linda.account /data/account**.
3. Set the permissions to enable the user and group owners to write files to these directories, and deny all access for all others: **chmod 770 /data/sales**, and next **chmod 770 /data/account**.
4. Use **su - laura** to become laura and change into the directory /data/account. Use **touch emptyfile** to create a file in this directory. Does this work?
5. Still as laura, use **cd /data/sales** and use **touch emptyfile** to create a file in this directory. Does this work?

Managing Advanced Permissions

Apart from the basic permissions that you have just read about, Linux has a set of advanced permissions as well. These are not permissions that you would set by default, but they sometimes provide a useful addition to realize more advanced scenarios. In this section, you learn what they are and how to set them.

Understanding Advanced Permissions

There are three advanced permissions. The first of them is the set user ID (SUID) permission. On some very specific occasions, you may want to apply this permission to executable files. By default, a user who runs an executable file runs this file with his

own permissions. For normal users, that usually means that the use of the program is restricted. In some cases, however, the user needs special permissions, just for the execution of a certain task.

Consider, for example, the situation where a user needs to change his password. To do this, the user needs to write his new password to the `/etc/shadow` file. This file, however, is not writeable for users who do not have root permissions:

```
[root@hnl ~]# ls -l /etc/shadow
-----. 1 root root 1184 Apr 30 16:54 /etc/shadow
```

The SUID permission offers a solution for this problem. On the `/usr/bin/passwd` utility, this permission is applied by default. That means that when changing his password, the user temporarily has root permissions, which allows him to write to the `/etc/shadow` file. You can see the SUID permission with `ls -l` as an `s` at the position where normally you would expect to see the `x` for the user permissions:

```
[root@hnl ~]# ls -l /usr/bin/passwd
-rwsr-xr-x. 1 root root 32680 Jan 28 2010 /usr/bin/passwd
```

The SUID permission may look useful (and it is in some cases), but at the same time, it is potentially dangerous. If applied wrongly, you may give away root permissions by accident. I therefore recommend using it with greatest care only, or better yet: don't use it at all. Most administrators will never have to use it; you'll only see it on some files where the operating system needs to set it as a default.

The second special permission is set group ID (SGID). This permission has two effects. If applied on an executable file, it gives the user who executes the file the permissions of the group owner of that file. So, SGID can accomplish more or less the same thing that does SUID. For this purpose, however, SGID is hardly used. As is the case for the SUID permission, SGID is applied to some system files as a default setting.

When applied to a directory, SGID may be useful, because you can use it to set default group ownership on files and subdirectories created in that directory. By default, when a user creates a file, his effective primary group is set as the group owner for that file. That is not always very useful, especially because on Red Hat Enterprise Linux, users have their primary group set to a group with the same name as the user, and of which the user is the only member. So by default, files that a user creates will be group shared with nobody else.

Imagine a situation where users `linda` and `lori` work for the accounting department and are both members of the group `account`. By default, these users are members of the private group of which they are the only member. Both users, however, are members of the accounting group as well but as a secondary group setting.

The default situation is that when either of these users creates a file, the primary group becomes owner. So by default, linda cannot access the files that lori has created and vice versa. However, if you create a shared group directory (say, /groups/account) and make sure that the SGID permission is applied to that directory, and that the group accounting is set as the group owner for that directory, all files created in this directory and all its subdirectories also get the group accounting as the default group owner. For that reason, the SGID permission is a very useful permission to set on shared group directories.

The SGID permission shows in the output of **ls -l** as an **s** at the position where you normally find the group execute permission:

```
[root@hnl data]# ls -ld account
drwxr-sr-x. 2 root account 4096 Apr 30 21:28 account
```

The third of the special permissions is sticky bit. This permission is useful to protect files against accidental deletion in an environment where multiple users have write permissions in the same directory. If sticky bit is applied, a user may delete a file only if she is the user owner of the file or of the directory that contains the file. It is for that reason applied as a default permission to the /tmp directory, and it can be useful on shared group directories as well.

Without sticky bit, if a user can create files in a directory, she can also delete files from that directory. In a shared group environment, this may be annoying. Imagine users linda and lori again, who both have write permissions to the directory /data/account and get these permissions because of their membership in the group accounting. Therefore, linda can delete files that lori has created and vice versa.

When you apply sticky bit, a user can delete files only if either of the following is true:

- The user is owner of the file.
- The user is owner of the directory where the file exists.

When using **ls -l**, you can see sticky bit as a **t** at the position where you normally see the execute permission for others:

```
[root@hnl data]# ls -ld account/
drwxr-sr-t. 2 root account 4096 Apr 30 21:28 account/
```

TIP Make sure that you know how to manage these advanced permissions. The RHCSA objectives specifically mention that you need to be able to use set GID to create a shared group directory.

Applying Advanced Permissions

To apply SUID, SGID, and sticky bit, you can use **chmod** as well. SUID has numeric value 4, SGID has numeric value 2, and sticky bit has numerical value 1. If you want to apply these permissions, you need to add a four-digit argument to **chmod**, of which the first digit refers to the special permissions. The following line would, for example, add the SGID permission to a directory, and set rwx for user and rx for group and others:

```
chmod 2755 /somedir
```

It is rather impractical if you have to look up the current permissions that are set before working with **chmod** in absolute mode. (You risk overwriting permissions if you do not.) Therefore, I recommend working in relative mode if you need to apply any of the special permissions:

- For SUID, use **chmod u+s**.
- For SGID, use **chmod g+s**.
- For sticky bit, use **chmod +t**, followed by the name of the file or the directory that you want to set the permissions on.

Table 7-4 summarizes all that is important to know about managing special permissions.

Table 7-4 Working with SUID, SGID, and Sticky Bit

Permission	Numeric Value	Relative Value	On Files	On Directories
SUID	4	u+s	User executes file with permissions of file owner.	No meaning.
SGID	2	g+s	User executes file with permissions of group owner.	Files created in directory get the same group owner.
Sticky bit	1	+t	No meaning.	Prevents users from deleting files from other users.

**Key
Topic**

In Exercise 7-2, you use special permissions to make it easier for members of a group to share files in a shared group directory. You assign the set group ID bit and sticky bit and see that after setting these, features are added that make it easier for group members to work together.

Exercise 7-2 Working with Special Permissions

1. Open a terminal in which you are user **linda**.
2. Use **cd /data/sales** to go to the sales directory. Use **touch linda1** and **touch linda2** to create two files of which **linda** is the owner.
3. Type **exit** to go back to a root shell, and next use **su - laura** to switch the current user identity to user **laura**, who also is a member of the sales group.
4. Use **cd /data/sales** again, and from that directory, use **ls -l**. You'll see the two files that were created by user **linda** that are group-owned by the group **linda**. Use **rm -f linda***. This will remove both files.
5. Use the commands **touch laura1 laura2** to create two files that are owned by user **laura**.
6. Use **su -** to escalate your current permissions to root level.
7. Use **chmod g+s,o+t /data/sales** to set the group ID bit as well as sticky bit on the shared group directory.
8. Use **su - linda** and type **cd /data/sales**. First, use **touch linda3 linda4**. You should now see that the two files you have created are owned by the group **sales**, which is group owner of the directory **/data/sales**.
9. Use **rm -rd laura***. Normally, sticky bit prevents you from doing so, but because **linda** is the owner of the directory that contains the files, you are allowed to do it anyway!

Managing ACLs

The advanced permissions that were discussed in the previous section add useful functionality to the way Linux works with permissions, but that functionality does not allow you to give permissions to more than one user or one group on the same file. Access control lists do offer this feature. Apart from that, they allow administrators to set default permissions in a sophisticated way.

Understanding ACLs

Although the ACL subsystem adds great functionality to your server, there is one drawback: not all utilities support it. Therefore, you might lose ACL settings when copying or moving files, and your backup software might not be able to back up ACL settings.

Previous versions of the **tar** utility did not support working with ACLs, which is why the **star** utility was introduced. You won't need this utility anymore because current versions of **tar** do support working with ACLs. The **star** utility, however, is still listed in the exam objectives.

TIP You can also use **getfacl** to create a backup of ACLs, which can be restored using the **setfacl** command. To create the backup, use **getfacl -R /directory > file.acls**. To restore the settings from the backup file, use **setfacl --restore=file.acl**.

The lack of support by some tools does not have to be a problem though. ACLs are often applied to directories as a structural measure and not on individual files. Therefore, you will not have lots of ACLs, but just a few applied in smart places in the file system. Hence, it is relatively easy to restore the original ACLs you were working with, even if your backup software does not support them.

Preparing Your File System for ACLs

Before starting to work with ACLs, you might have to prepare your file system for ACL support. Because the file system metadata needs to be extended, there is not always default support for ACLs in the file system. If while setting ACLs to a file system you are getting an “operation not supported” message, your file system probably lacks support for ACLs. To fix this, you need to add the **acl** mount option in the **/etc/fstab** file so that the file system will be mounted with ACL support by default. Read Chapter 14, “Managing Storage,” for more information about the **/etc/fstab** file and options that you can use in that file.

Changing and Viewing ACL Settings with **setfacl** and **getfacl**

To set ACLs, you use the **setfacl** command. To see current ACL settings, you use **getfacl**. The **ls -l** command does not show any existing ACLs; it just shows a **+** after the listing of the permissions, which indicates that ACLs apply to the file as well. Before setting ACLs, it is always a good idea to show current ACL settings using **getfacl**. In Example 7-3, you can see the current permissions as shown with **ls -l** and also as shown with **getfacl**. If you look closely, you can see that the information shown is exactly the same.

Example 7-3 Checking Permissions with **ls -l** and **getfacl**

```
[root@server1 /]# ls -ld /dir
drwxr-xr-x. 2 root root 6 Feb 6 11:28 /dir
[root@server1 /]# getfacl /dir
getfacl: Removing leading '/' from absolute path names
# file: dir
# owner: root
# group: root
user::rwx
group::r-x
other::r-x
```

In the result of the **getfacl** command in Example 7-3, you can see that the permissions are shown for three different entities: the user, the group, and others. Now let's add an ACL to give read and execute permissions to the group sales as well. The command to use for this is **setfacl -m g:sales:rx /dir**. In this command, **-m** indicates that the current ACL settings need to be modified. After that, **g:sales:rx** tells the command to set the ACL to read and execute (**rx**) for the group (**g**) sales. In Example 7-4, you can see what the command looks like, as well as the output of the **getfacl** command after changing the current ACL settings.

Example 7-4 Changing Group ACLs Using **setfacl**

```
[root@server1 /]# setfacl -m g:sales:rx /dir
[root@server1 /]# getfacl /dir
getfacl: Removing leading '/' from absolute path names
# file: dir
# owner: root
# group: root
user::rwx
group::r-x
group:sales:r-x
mask::r-x
other::r-x
```

Now that you understand how to set a group ACL, it is easy to understand ACLs for users and others as well. For instance, the command **setfacl -m u:linda:rwx /data** gives permissions to user linda on the /data directory without making her the owner and without changing the current owner assignment.

The **setfacl** command has many possibilities and options. One option is particularly important: the option **-R**. If used, the option applies the ACL to all files and subdirectories currently existing in the directory where you set the ACL. It is a good idea to always use this option while changing ACLs for existing directories.

ACLs also allow you to take away permissions from users. For instance, **setfacl -m u:anna:- /tmp/myfile** would take away all permissions for user anna. This also will overwrite any permissions that the user may have obtained through the others entity.

Working with Default ACLs

One benefit of using ACLs is that you can give permissions to more than one user or group at a directory. Another benefit is that you can enable inheritance by working with default ACLs. By setting a default ACL, you determine the permissions that will be set for all new items that are created in the directory. Be aware, though, that a default ACL does not change the permissions for existing files and subdirectories. To change those as well you need to add a normal ACL, too!

TIP This is important to know. If you want to use ACLs to configure access for multiple users or groups to the same directory, you have to set ACLs twice. First, use **setfacl -R -m** to modify the ACLs for current files. Then, use **setfacl -m d:** to take care of all new items that will be created also.

To set a default ACL, simply add the option **d** after the option **-m**. (Order does matter!) So, use **setfacl -m d:g:sales:rx /data** if you want group sales to have read and execute permissions on everything that will ever be created in the /data directory.

When using default ACLs, it can also be useful to set an ACL for the others entity. Normally, this does not make much sense because you can also change the permissions for others by using **chmod**. What you cannot do using **chmod**, though, is specify the permissions that should be given to others on every new file that will ever be created. If you want others not to get permissions on anything that is created in /data, for example, use **setfacl -m d:o::- /data**.

In Exercise 7-3, you learn how to apply advanced permissions and ACLs by working further on the /data/account and /data/sales directories that you previously created. In earlier exercises, you ensured that the group sales has permissions on /data/sales, and the group account has permissions on /data/account. Up to now, you have not been able to upgrade the design, which you do in this exercise. You first make sure that the group account gets read permissions on the /data/sales directory and that group sales gets read permissions on the /data/account directory. Then, you set default ACLs to make sure that on all new files the permissions are properly set for all new items.

TIP ACLs and regular permissions are not always very well integrated. Problems may arise if you have applied default ACLs to a directory, after which items have been added to that directory, and then you try to change the regular permissions. The changes that are applied to the regular permissions will not be reflected well in the ACL overview. To avoid problems, set regular permissions first, after which you set default ACLs (and after doing that, try to not change them again).

Exercise 7-3 Managing Advanced Permissions Using ACLs

1. Open a root terminal.
2. Use **setfacl -m g:account:rx /data/sales** to give the group account read permissions on the /data/sales directory, and use **setfacl -m g:sales:rx /data/account** to give the group sales read permissions on the /data/account directory.
3. Use **getfacl /data/** to verify that the permissions have been set the way you intended to.
4. Use **setfacl -m d:g:account:rx,g:sales:rwX /data/sales** to set the default ACL for the directory sales.
5. Add the default ACL for the directory /data/account by using **setfacl -m d:g:sales:rx,g:account:rwX /data/account**.
6. Verify that the ACL settings are effective by adding a new file in /data/sales. Use **touch /data/sales/newfile** and use **getfacl /data/sales/newfile** to check the current permission assignments.

Setting Default Permissions with umask

In the previous section about ACLs, you learned how to work with default ACLs. If you do not use ACLs, there is a shell setting that determines the default permissions that you will get: the **umask**. In this section, you learn how to modify default permissions using **umask**.

You have probably noticed that when creating a new file, some default permissions are set. These permissions are determined by the umask setting. This shell setting is applied to all users when logging in to the system. In the umask setting, a numeric value is used that is subtracted from the maximum permissions that can be set automatically to a file; the maximum setting for files is 666, and for directories is 777. Some exceptions apply to this rule, however. You can find a complete overview of umask settings in Table 7-5.

Of the digits used in the umask, like with the numeric arguments for the **chmod** command, the first digit refers to user permissions, the second digit refers to the group permissions, and the last refers to default permissions set for others. The default umask setting of 022 gives 644 for all new files and 755 for all new directories that are created on your server. A complete overview of all umask numeric values and their result is shown in Table 7-5.

Table 7-5 umask Values and Their Result

Value	Applied to Files	Applied to Directories
0	Read and write	Everything
1	Read and write	Read and write
2	Read	Read and execute
3	Read	Read
4	Write	Write and execute
5	Write	Write
6	Nothing	Execute
7	Nothing	Nothing

**Key
Topic**

An easy way to see how the umask setting works is as follows: Start with the default permissions for a file set to 666 and subtract the umask to get the effective permissions. For a directory, start with its default permissions that are set to 777 and subtract the umask to get the effective permissions.

There are two ways to change the umask setting: for all users and for individual users. If you want to set the umask for all users, you must make sure the umask setting is considered when starting the shell environment files as directed by `/etc/profile`. The right approach is to create a shell script with the name `umask.sh` in the `/etc/profile.d` directory and specify the umask you want to use in that shell script. If the umask is changed in this file, it applies to all users after logging in to your server.

An alternative to setting the umask through `/etc/profile` and related files where it is applied to all users logging in to the system is to change the umask settings in a file with the name `.profile`, which is created in the home directory of an individual user. Settings applied in this file are applied for the individual user only; therefore, this is a nice method if you need more granularity. I personally like this feature to change the default umask for user root to 027, whereas normal users work with the default umask 022.

Working with User-Extended Attributes

When you work with permissions, a relationship always exists between a user or group object and the permissions these user or group objects have on a file or directory. An alternative method of securing files on a Linux server is by working with attributes. Attributes do their work regardless of the user who accesses the file.

As is the case for ACLs, for file attributes a **mount** option, **user_xattr**, may have to be enabled. If you are getting an “operation not supported” message while working with user-extended attributes, make sure to set this **mount** option in the `/etc/fstab` file.

Many attributes are documented. Some attributes are available but not yet implemented. Do not use them; they bring you nothing. Following are the most useful attributes that you can apply:

- **A** This attribute ensures that the file access time of the file is not modified. Normally, every time a file is opened, the file access time must be written to the file’s metadata. This affects performance in a negative way; therefore, on files that are accessed on a regular basis, the **A** attribute can be used to disable this feature.
- **a** This attribute allows a file to be added to but not to be removed.
- **c** If you are using a file system where volume level compression is supported, this file attribute makes sure that the file is compressed the first time the compression engine becomes active.
- **D** This attribute makes sure that changes to files are written to disk immediately, and not to cache first. This is a useful attribute on important database files to make sure that they do not get lost between file cache and hard disk.
- **d** This attribute makes sure the file is not backed up in backups where the **dump** utility is used.
- **I** This attribute enables indexing for the directory where it is enabled. This allows faster file access for primitive file systems like Ext3 that do not use a B-tree database for fast access to files.
- **i** This attribute makes the file immutable. Therefore, no changes can be made to the file at all, which is useful for files that need a bit of extra protection.
- **j** This attribute ensures that on an Ext3 file system the file is first written to the journal and only after that to the data blocks on the hard disk.

- **s** This attribute overwrites the blocks where the file was stored with 0s after the file has been deleted. This makes sure that recovery of the file is not possible after it has been deleted.
- **u** This attribute saves undelete information. This allows a utility to be developed that works with that information to salvage deleted files.

NOTE Although there are quite a few attributes that can be used, be aware that most attributes are rather experimental and are only of any use if an application is used that can work with the given attribute. For example, it does not make sense to apply the **u** attribute if no application has been developed that can use this attribute to recover deleted files.

If you want to apply attributes, you can use the **chattr** command. For example, use **chattr +s somefile** to apply the attributes to somefile. Need to remove the attribute again? Then use **chattr -s somefile** and it will be removed. To get an overview of all attributes that are currently applied, use the **lsattr** command.

Summary

In this chapter, you learned how to work with permissions. You read about the three basic permissions, the advanced permissions, and how to apply ACLs on the file system. You also learned how to use the umask setting to apply default permissions. Toward the end of this chapter, you learned how to use user-extended attributes to apply an additional level of file system security.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 7-6 lists a reference of these key topics and the page number on which each is found.

**Key
Topic****Table 7-6** Key Topics for Chapter 7

Key Topic Element	Description	Page
Table 7-2	Use of the basic permissions	150
Table 7-3	Numeric representation of permissions	151
Table 7-4	Special permissions overview	155
Table 7-5	umask overview	161

Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

ownership, permissions, access control list, inheritance, attribute

Review Questions

The questions that follow use an open-ended format that is meant to help you test your knowledge of the concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. How do you use **chown** to set the group owner to a file?
2. Which command finds all files that are owned by a specific user?
3. How would you apply read, write, and execute permissions to all files in /data for the user and group owners while setting no permissions for others?
4. Which command enables you in relative permission mode to add the execute permission to a file that you want to make executable?
5. Which command enables you to ensure that group ownership of all new files that will be created in a directory is set to the group owner of that directory?
6. You want to ensure that users can only delete files of which they are the owner, or files that are in a directory of which they are the owner. Which command will do that for you?

7. Which command adds an ACL that grants members of the group sales read permissions to all existing files in the current directory?
8. What do you need to do to ensure that members of the group sales get read permissions to all files in the current directory and all of its subdirectories, as well as all files that will be created in this directory in the future?
9. Which umask do you need to set if you never want “others” to get any permissions on new files?
10. Which command ensures that nobody can delete myfile by accident?

End-of-Chapter Lab

In Chapter 6, “User and Group Management,” you created some users and groups. These users and groups are needed to perform the exercises in this lab.

Lab 7.1

1. Set up a shared group environment. Create two directories: /data/account and /data/sales. Make the group sales the owner of the directory sales, and make the group account the owner of the directory account.
2. Configure the permissions so that the user owner (which must be root) and group owner have full access to the directory. There should be no permissions assigned to the others entity. Also make sure that “others” will get no permissions on newly created files and directories within the entire /data structure.
3. Create a configuration that allows members of the group sales to read files in the directory /data/account. Also make sure that members of the group account have read permissions in the directory /data/sales.
4. Ensure that all new files in both directories inherit the group owner of their respective directory. This means that all files that will be created in /data/sales will be owned by the group sales, and all files in /data/account will be owned by the group account.
5. Ensure that users are only allowed to remove files of which they are the owner.



This chapter covers the following subjects:

- Networking Fundamentals
- Managing Network Addresses and Interfaces
- Validating Network Configuration
- Managing Network Configuration with **nmtui** and **nmcli**

The following RHCSA exam objectives are covered:

- Configure IPv4 and IPv6 addresses
- Configure hostname resolution

Networking is one of the most essential items on a modern server. On RHEL 8, networking is managed by the NetworkManager service, and with the release of RHEL 7, some completely new tools were introduced to help manage networks. If you have already worked with networking on RHEL 6 and older, you will notice that networking has changed a lot and that some approaches that were the default in earlier versions will no longer work.

Configuring Networking

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 8-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 8-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Networking Fundamentals	1–2
Managing Network Addresses and Interfaces	3
Validating Network Configuration	4
Configuring Network Configuration with nmtui and nmcli	5–8
Setting Up Hostname and Name Resolution	9–10

- Which of the following IP addresses belong to the same network?
 - 192.168.4.17/26
 - 192.168.4.94/26
 - 192.168.4.97/26
 - 192.168.4.120/26
 - I and II
 - II and III
 - III and IV
 - II, III, and IV

2. Which of the following is *not* a private IP address?
 - a. 10.10.10.10
 - b. 169.254.11.23
 - c. 172.19.18.17
 - d. 192.168.192.192
3. Which of the following would be the default network interface name on a RHEL 8 system?
 - a. p6p1
 - b. eth0
 - c. eno1677783
 - d. e0
4. Which command shows the recommended way to display information about the network interface as well as its IP configuration?
 - a. **ifconfig -all**
 - b. **ipconfig**
 - c. **ip link show**
 - d. **ip addr show**
5. Which statement about NetworkManager is *not* true?
 - a. It is safe to disable NetworkManager and work with the network service instead.
 - b. NetworkManager manages network connections that are applied to network interfaces.
 - c. NetworkManager has a text-based user interface with the name **nmtui**.
 - d. NetworkManager is the default service to manage networking in RHEL 8.
6. Which man page contains excellent examples on **nmcli** usage?
 - a. nmcli
 - b. nmcli-examples
 - c. nm-config
 - d. nm-tools

7. Which of the following is the name of the text user interface to specify network connection properties?
 - a. system-config-network
 - b. system-config-networkmanager
 - c. nmtui
 - d. nmcli
8. Which of the following commands shows correct syntax to set a fixed IP address to a connection using **nmcli**?
 - a. **nmcli con add con-name "static" ifname eth0 autoconnect no type ethernet ipv4 10.0.0.10/24 gw4 10.0.0.1**
 - b. **nmcli con add con-name "static" ifname eth0 autoconnect no type ethernet ipv4 10.0.0.10/24 gwv4 10.0.0.1**
 - c. **nmcli con add con-name "static" ifname eth0 type ethernet ipv4 10.0.0.10/24 gw4 10.0.0.1**
 - d. **nmcli con add con-name "static" ifname eth0 autoconnect no type ethernet ip4 10.0.0.10/24 gw4 10.0.0.1**
9. Which of the following is *not* a recommended way to specify which DNS servers to use?
 - a. Edit `/etc/resolv.conf`.
 - b. Set the DNS options in `/etc/sysconfig/network-scripts/ifcfg-<ID>`.
 - c. Set the DNS server names using **nmcli**.
 - d. Use **nmtui** to set the DNS server names.
10. In which configuration file would you set the hostname?
 - a. `/etc/sysconfig/network`
 - b. `/etc/sysconfig/hostname`
 - c. `/etc/hostname`
 - d. `/etc/defaults/hostname`

Foundation Topics

Networking Fundamentals

To set up networking on a server, your server needs a unique address on the network. For this purpose, IP (Internet Protocol) addresses are used. Currently, two versions of IP addresses are relevant:

Key Topic

- **IPv4 addresses:** These are based on 32-bit addresses and have four octets, separated by dots, such as 192.168.10.100.
- **IPv6 addresses:** These are based on 128-bit addresses and are written in eight groups of hexadecimal numbers that are 16 bits each and separated by colons. An IPv6 address may look like fe80:badb:abe01:45bc:34ad:1313:6723:8798.

In this chapter, you learn how to work with IPv4 addresses. IPv6 addresses are described only briefly.

IP Addresses

Originally, IP addresses were assigned to computers and routers. Nowadays, many other devices also need IP addresses to communicate, such as smartphones, industrial equipment, and almost all other devices that are connected to the Internet. This chapter refers to all of those devices by using the word *node*. You'll also occasionally encounter the word *host*. A host is typically a server providing services on the network.

To make it easier for computers to communicate with one another, every IP address belongs to a specific network, and to communicate with computers on another network, a router is used. A router is a machine (often dedicated hardware that has been created for that purpose) that connects networks to one another.

To communicate on the Internet, every computer needs a worldwide unique IP address. These addresses are scarce; a theoretical maximum of four billion IP addresses is available, and that is not enough to provide every device on the planet with an IP address. IPv6 is the ultimate solution for that problem, because a very large number of IP addresses can be created in IPv6. Because many networks still work with IPv4, though, another solution exists: the private network addresses.

Private network addresses are addresses that are for use in internal networks only. Some specific IP network addresses have been reserved for this purpose:

Key Topic

- 10.0.0.0/8 (a single Class A network)
- 172.16.0.0/12 (16 Class B networks)
- 192.168.0.0/16 (256 Class C networks)

When private addresses are used, the nodes that are using them cannot access the Internet directly, and nodes from the Internet cannot easily access them. Because that is not very convenient, Network Address Translation (NAT) is commonly used on the router that connects the private network to the Internet. In NAT, the nodes use a private IP address, but when accessing the Internet, this private IP address is replaced with the IP address of the NAT router. Hence, nodes on the Internet think that they are communicating with the NAT router, and not with the individual hosts.

The NAT router on its turn uses tables to keep track of all connections that currently exist for the hosts in the network. Based on this table, the NAT router helps make it possible for computers with a private IP address to connect to hosts on the Internet anyway. The use of NAT is very common; it is embedded in most routers that are used in home and small business networks to connect computers and other devices in those networks to the Internet.

IPv6 Addresses

Let's have a look at a valid IPv6 address, such as 02fb:0000:0000:0000:90ff:fe23:8998:1234. In this address, you can see that a long range of zeroes occurs. To make IPv6 addresses more readable, you can replace one range of zeros with ::. Also, if an IPv6 address starts with a leading zero, you can omit it. So the previously mentioned IPv6 address can be rewritten as 2fb::90ff:fe23:8998:1234.

Network Masks

To know to which network a computer belongs, a subnet mask is used with every IP address. The *subnet mask* defines which part of the network address indicates the network and which part indicates the node. Subnet masks may be written in the classless interdomain routing (CIDR) notation, which indicates the number of bits in the subnet mask, or in the classical notation, and they always need to be specified with the network address. Examples include 192.168.10.100/24 (CIDR notation), which indicates that a 24-bit network address is used, and 192.168.10.100/255.255.255.0 (classical notation), which indicates exactly the same.

Often, network masks use multiple bytes. In the example using 192.168.10.100/24, the first three bytes (the 192.168.10 part) form the network part, and the last byte (the number 100) is the host part on that network.

When talking about network addresses, you use a 4-byte number, as well, in which the node address is set to 0. So in the example of 192.168.10.100/24, the network address is 192.168.10.0. In IPv4 networks, there is also always a *broadcast address*. This is the address that can be used to address all nodes in the network. In the broadcast address, all node bits are set to 1, which makes for the decimal number 255 if an entire byte is referred to. So in the example of the address 192.168.10.100/24, the broadcast address is 192.168.10.255.

Binary Notation

Because the number of IPv4 addresses is limited, in modern IPv4 networks variable-length network masks are used. These are network masks such as 212.209.113.33/27. In a variable-length subnet mask, only a part of the byte is used for addressing nodes, and another part is used for addressing the network. In the subnet mask /27, the first 3 bits of the last byte are used to address the network, and the last 5 bits are used for addressing nodes. This becomes clearer if you write down the address in a binary notation:

IP address:

212.209.113.33 = 11010100.11010001.00001010.00100001

Subnet mask:

/27 = 11111111.11111111.11111111.11100000

When applying the subnet mask to the IP address, you can see that the first 3 bits of the IP address belong to the network, so the network is 00100000. And if you use a binary calculator, you can see that that corresponds with the decimal IP address 32. Using the /27 subnet mask allows for the creation of multiple networks. Table 8-2 gives an overview.



Table 8-2 Binary-Decimal Conversion Overview

Binary Value	Decimal Value
00000000	0
00100000	32
01000000	64
01100000	96
10000000	128
10100000	160
11000000	192
11100000	224

So, based on this information, if you consider the IP address 212.209.113.33/27 again, you can see that it belongs to the network 212.209.113.32/27, and that in this network the broadcast address (which has the node part of the IP address set to all 1s) is 212.209.113.63; therefore, with a /27 subnet mask, 30 nodes can be addressed per network. You'll get 32 IP addresses, but 2 of them are the network address and the broadcast address, which cannot be used as a host IP address.

EXAM TIP You do not need to make this kind of calculation on the RHCSA exam, but it helps understanding how IP network addressing works.

MAC Addresses

IP addresses are the addresses that allow nodes to communicate to any other node on the Internet. They are not the only addresses in use though. Each network card also has a 12-byte MAC address. MAC addresses are for use on the local network (that is, the local physical network or local WLAN, just up to the first router that is encountered); they cannot be used for communications between nodes that are on different networks. MAC addresses are important, though, because they help computers find the specific network card that an IP address belongs to.

An example of a MAC address is 00:0c:29:7d:9b:17. Notice that each MAC address consists of two parts. The first 6 bytes is the vendor ID, and the second 6 bytes is the unique node ID. Vendor IDs are registered, and by using registered vendor IDs, it is possible to allocate unique MAC addresses.

Protocol and Ports

In the previous section you learned how to use IP addresses to identify individual nodes. On these nodes, you will typically be running services, like a web server or an FTP server. To identify these services, port addresses are used. Every service has a specific port address, such as port 80 for Hypertext Transfer Protocol (HTTP) or port 22 for a Secure Shell (SSH) server, and in network communication, the sender and the receiver are using port addresses. So, there is a destination port address as well as a source port address involved in network communications.

Because not all services are addressed in a similar way, a specific protocol is used between the IP address and the port address, such as Transfer Control Protocol (TCP), User Datagram Protocol (UDP), or Internet Control Message Protocol (ICMP). Every protocol has specific properties: TCP is typically used when the network communication must be reliable and delivery must be guaranteed; UDP is used when it must be fast and guaranteed delivery is not necessary.

Managing Network Addresses and Interfaces

As a Linux server administrator, you need to manage network addresses and network interfaces. The network addresses can be assigned in two ways:

Key Topic

- **Fixed IP addresses:** Useful for servers that always need to be available at the same IP address.
- **Dynamically assigned IP addresses:** Useful for end users' devices, and for instances in a cloud environment. To dynamically assign IP addresses, a Dynamic Host Configuration Protocol (DHCP) server is usually used.

For a long time, network cards in Linux have had default names, such as `eth0`, `eth1`, and `eth2`. This naming was assigned based on the order of detection of the network card. So, `eth0` was the first network card that got detected, `eth1` the second, and so on. This worked well in an environment where a node has one or two network cards only. If a node has multiple network cards that need to be dynamically added and removed, however, this approach does not work so well because it is very hard to identify which physical network card is using which name.

In RHEL 8, the default names for network cards are based on firmware, device topology, and device types. This leads to network card names that always consist of the following parts:

- Ethernet interfaces begin with *en*, WLAN interfaces begin with *wl*, and WWAN interfaces begin with *ww*.
- The next part of the name represents the type of adapter. An *o* is used for onboard, *s* is for a hotplug slot, and *p* is for a PCI location. Administrators can also use *x* to create a device name that is based on the MAC address of the network card.
- Then follows a number, which is used to represent an index, ID, or port.
- If the fixed name cannot be determined, traditional names such as `eth0` are used.

Based on this information, device names such as `eno16777734` can be used, which stands for an onboard Ethernet device, with its unique index number.

Apart from this default device naming, network cards can be named based on the BIOS device name as well. In this naming scheme, names such as `em1` (embedded network card 1) or `p4p1` (which is PCI slot 4, port 1) can be used. To use this kind of naming, the `biosdevname` package must be installed.

Validating Network Configuration

Before you can learn how to set network information, you must know how to verify current network information. In this section, you learn how to do that, and you learn how to check the following networking items:

- IP address and subnet mask
- Routing
- Availability of ports and services

Validating Network Address Configuration

To verify the configuration of the network address, you need to use the **ip** utility. The **ip** utility is a modern utility that can consider advanced networking features that have been introduced recently. With the **ip** utility, many aspects of networking can be monitored:

- Use **ip addr** to configure and monitor network addresses.
- Use **ip route** to configure and monitor routing information.
- Use **ip link** to configure and monitor network link state.

Apart from these items, the **ip** utility can manage many other aspects of networking, but you do not need to know about them for the RHCSA exam.

WARNING In earlier Linux versions and some other UNIX-like operating systems, the **ifconfig** utility was and is used for validating network configuration. Do not use this utility on modern Linux distributions. Because Linux has become an important player in cloud computing, networking has evolved a lot to match cloud computing requirements, and many new features have been added to Linux networking. With the **ifconfig** utility, you cannot manage or validate these concepts. Even if **ifconfig** is still the default tool on some operating systems (like Apple OS X for instance), you should never use it anymore on Linux!

To show current network settings, you can use the **ip addr show** command (which can be abbreviated as **ip a s** or even as **ip a**). The **ip** command is relatively smart and does not always require you to type the complete option.

The result of the **ip addr show** command looks like Example 8-1.

Example 8-1 Monitoring Current Network Configuration with **ip addr show**

```
[root@server1 ~]# ip addr show
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
    group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: ens33: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel
    state UP group default qlen 1000
    link/ether 00:0c:29:50:9e:c9 brd ff:ff:ff:ff:ff:ff
    inet 192.168.4.210/24 brd 192.168.4.255 scope global dynamic
        noprefixroute ens33
        valid_lft 1370sec preferred_lft 1370sec
    inet6 fe80::959:3b1a:9607:8928/64 scope link noprefixroute
        valid_lft forever preferred_lft forever
```

In the result of this command, you see a listing of all network interfaces in your computer. You'll normally see at least two interfaces, but on specific configurations, there can be many more interfaces. In Example 8-1, two interfaces are shown: the loopback interface `lo`, and the onboard Ethernet card `ens33`.

The loopback interface is used for communication between processes. Some processes are using the IP protocol for internal communications. For that reason, you'll always find a loopback interface, and the IP address of the loopback interface is always set to 127.0.0.1. The important part of the output of the command is for the onboard Ethernet card. The command shows the following items about its current status:

- **Current state:** The most important part of this line is the text `state UP`, which shows that this network card is currently up and available.
- **MAC address configuration:** This is the unique MAC address that is set for every network card. You can see the MAC address itself (`00:0c:29:50:9e:c9`), as well as the corresponding broadcast address.
- **IPv4 configuration:** This line shows the IP address that is currently set, as well as the subnet mask that is used. You can also see the broadcast address that is used for this network configuration. Notice that on some interfaces you may find multiple IPv4 addresses.
- **IPv6 configuration:** This line shows the current IPv6 address and its configuration. Even if you haven't configured anything, every interface automatically gets an IPv6 address, which can be used for communication on the local network only.

If you are just interested in the link state of the network interfaces, you can use the **ip link show** command. This command (of which you can see the output in Example 8-2) repeats the link state information of the **ip addr show** command. If you add the option **-s**, you can also see current link statistics, which gives information about packets transmitted and received, as well as an overview of errors that have occurred during packet transmission.

Example 8-2 ip link show Output

```
[root@server1 ~]# ip -s link show
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
   mode DEFAULT group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    RX: bytes    packets  errors  dropped overrun mcast
         0         0        0       0        0        0
    TX: bytes    packets  errors  dropped carrier collsns
         0         0        0       0        0        0
2: ens33: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel
   state UP mode DEFAULT group default qlen 1000
    link/ether 00:0c:29:50:9e:c9 brd ff:ff:ff:ff:ff:ff
    RX: bytes    packets  errors  dropped overrun mcast
   143349       564        0       0        0        0
    TX: bytes    packets  errors  dropped carrier collsns
   133129       541        0       0        0        0
```

In case the **ip link show** command shows the current link state as down, you can temporarily bring it up again by using **ip link set**, which is followed by **dev devicename** and **up** (for example, **ip link set dev ens33 up**).

In Exercise 8-1, you learn how to manage and monitor networking with the **ip** utility and other utilities.

Exercise 8-1 Validating Network Configuration

1. Open a root shell.
2. Type **ip -s link**. This shows all existing network connections, in addition to statistics about the number of packets that have been sent and associated error messages.
3. Type **ip addr show**. You'll see the current address assignments for network interfaces on your server.

Validating Routing

One important aspect of networking is routing. On every network that needs to communicate to nodes on other networks, routing is a requirement. Every network has, at least, a default router (also called the default gateway) that is set, and you can see which router is used as the default router by using the command **ip route show** (see Example 8-3). You should always perform one quick check to verify that your router is set correctly: the default router at all times must be on the same network as the local IP address that your network card is using.

Example 8-3 ip route show Output

```
[root@server1 ~]# ip route show
default via 192.168.4.2 dev ens33 proto dhcp metric 100
192.168.4.0/24 dev ens33 proto kernel scope link src 192.168.4.210
metric 100
192.168.122.0/24 dev virbr0 proto kernel scope link src 192.168.122.1
linkdown
```

In Example 8-3, the most important part is the first line. It shows that the default route goes through (“via”) IP address 192.168.4.2, and also shows that network interface ens33 must be used to address that IP address. The line shows that this default route was assigned by a DHCP server. The metric is used in case multiple routes are available to the same destination. In that case, the route with the lowest metric will be used. This is something important on router devices, but on computers that are not a router, the metric doesn’t really matter.

Next you can see lines that identify the local connected networks. When booting, an entry is added for each local network as well, and in this example this applies to the networks 192.168.4.0 and 192.168.122.0. These routes are automatically generated and do not need to be managed.

Validating the Availability of Ports and Services

Network problems can be related to the local IP address and router settings but can also be related to network ports that are not available on your server or on a remote server. To verify availability of ports on your server, you can use the **netstat** command, or the newer **ss** command, which provides the same functionality. Exercise 8-2 shows how to verify network settings. By typing **ss -lt**, you’ll see all listening TCP ports on the local system (see Example 8-4).

Example 8-4 Using **ss -lt** to Display All Listening Ports on the Local System

```
[root@server1 ~]# ss -lt
```

State	Recv-Q	Send-Q	Local Address:Port	Peer
LISTEN	0	32	192.168.122.1:domain	
0.0.0.0:*				
LISTEN	0	128	0.0.0.0:ssh	
0.0.0.0:*				
LISTEN	0	5	127.0.0.1:ipp	
0.0.0.0:*				
LISTEN	0	128	0.0.0.0:sunrpc	
0.0.0.0:*				
LISTEN	0	128	:::ssh	
:::*				
LISTEN	0	5	:::1:ipp	
:::*				
LISTEN	0	128	:::sunrpc	
:::*				

Notice where the port is listening on. Some ports are only listening on the IPv4 loopback address 127.0.0.1 or the IPv6 loopback address ::1, which means that they are locally accessible only and cannot be reached from external machines. Other ports are listening on *, which stands for all IPv4 addresses, or on :::*, which represents all ports on all IPv6 addresses.

Exercise 8-2 Verifying Network Settings

1. Open a root shell to your server and type **ip addr show**. This shows the current network configuration. Note the IPv4 address that is used and the network device names that are used; you need these later in this exercise.
2. Type **ip route show** to verify routing configuration.
3. If your computer is connected to the Internet, you can now use the **ping** command to verify the connection to the Internet is working properly. Type **ping -c 4 8.8.8.8**, for instance, to send four packets to IP address 8.8.8.8. If your Internet connection is up and running, you should get “echo reply” answers.
4. Type **ip addr add 10.0.0.10/24 dev <yourdevicename>**. This will temporarily set a new IP address.
5. Type **ip addr show**. You’ll see the newly set IP address, in addition to the IP address that was already in use.

6. Type **ifconfig**. Notice that you do not see the newly set IP address (and there are no options with the **ifconfig** command that allow you to see it). This is one example of why you should not use the **ifconfig** command anymore.
7. Type **ss -tul**. You'll now see a list of all UDP and TCP ports that are listening on your server.

Configuring Network Configuration with **nmtui** and **nmcli**

As mentioned earlier in this chapter, networking on RHEL 8 is managed by the NetworkManager service. You can use the **systemctl status NetworkManager** command to verify its current status. When NetworkManager comes up, it reads the network card configuration scripts, which are in `/etc/sysconfig/network-scripts` and have a name that starts with `ifcfg` and is followed by the name of the network card.

When working with network configuration in RHEL 8, you should know the difference between a device and a connection:

- A device is a network interface card.
- A connection is the configuration that is used on a device.


In RHEL 8, you can create multiple connections for a device. This makes sense on mobile computers, for example, to differentiate between settings that are used to connect to the home network and settings that are used to connect to the corporate network. Switching between connections on devices is something that is common on end-user computers, and not so common on servers. To manage the network connections that you want to assign to devices, you use the **nmtui** command or the **nmcli** command.

EXAM TIP Red Hat wants you to know how to work with **nmcli**. This command is not very easy to use, however, and in the end, on the exam you will need to configure a network device with the appropriate settings. For that reason, on the RHCSA exam, it is perfectly fine to use the **nmtui** text user interface command; you will get things done a lot easier with this command.

Required Permissions to Change Network Configuration

Obviously, the root user can make modifications to current networking. However, if an ordinary user is logged in to the local console, this user is able to make changes to the network configuration as well. As long as the user is using the

system keyboard to enter either a graphical console or a text-based console, these permissions are granted. This is because users are supposed to be able to connect their local system to a network. Notice that regular users that have used **ssh** to connect to a server are not allowed to change the network configuration. To check your current permissions, use the **nmcli gen permissions** command, as shown in Figure 8-1.



```
File Edit View Search Terminal Help
[student@server1 ~]$ nmcli gen permissions
PERMISSION                                     VALUE
org.freedesktop.NetworkManager.enable-disable-network    yes
org.freedesktop.NetworkManager.enable-disable-wifi       yes
org.freedesktop.NetworkManager.enable-disable-wwan       yes
org.freedesktop.NetworkManager.enable-disable-wimax      yes
org.freedesktop.NetworkManager.sleep-wake                no
org.freedesktop.NetworkManager.network-control           yes
org.freedesktop.NetworkManager.wifi.share.protected      yes
org.freedesktop.NetworkManager.wifi.share.open           yes
org.freedesktop.NetworkManager.settings.modify.system    yes
org.freedesktop.NetworkManager.settings.modify.own       yes
org.freedesktop.NetworkManager.settings.modify.hostname  auth
org.freedesktop.NetworkManager.settings.modify.global-dns auth
org.freedesktop.NetworkManager.reload                    auth
org.freedesktop.NetworkManager.checkpoint-rollback       auth
org.freedesktop.NetworkManager.enable-disable-statistics yes
org.freedesktop.NetworkManager.enable-disable-connectivity-check yes
[student@server1 ~]$
```


FIGURE 8-1 Verifying Current Permissions to Change Network Configuration

Configuring the Network with nmcli

Earlier in this chapter, you learned how to use **ip** to verify network configuration. You have also applied the **ip addr add** command to temporarily set an IP address on a network interface. Everything you do with the **ip** command, though, is nonpersistent. If you want to make your configuration persistent, use **nmtui** or **nmcli**.

A good start is to use **nmcli** to show all connections. This shows active *and* inactive connections. You can easily see the difference because inactive connections are not currently assigned to a device (see Example 8-5).

Example 8-5 Showing Current Connection Status



```
[root@server1 ~]# nmcli con show
```

NAME	UUID	TYPE	DEVICE
ens33	db6f53bd-654e-45dd-97ef-224514f8050a	ethernet	ens33

After finding the name of the connection, you can use **nmcli con show** followed by the name of the connection to see all properties of the connection. Notice that this command shows many properties. Example 8-6 shows the output of this command.

Example 8-6 Displaying Connection Properties

```
[root@server1 ~]# nmcli con show ens33
connection.id: ens33
connection.uuid: db6f53bd-654e-45dd-97ef-224514f8050a
connection.stable-id: --
connection.type: 802-3-ethernet
connection.interface-name: ens33
connection.autoconnect: yes
connection.autoconnect-priority: 0
connection.autoconnect-retries: -1 (default)
connection.multi-connect: 0 (default)
connection.auth-retries: -1
connection.timestamp: 1558778720
connection.read-only: no
connection.permissions: --
connection.zone: --
connection.master: --
connection.slave-type: --
connection.autoconnect-slaves: -1 (default)
connection.secondaries: --
connection.gateway-ping-timeout: 0
connection.metered: unknown
connection.lldp: default
connection.mdns: -1 (default)
connection.llmnr: -1 (default)
802-3-ethernet.port: --
802-3-ethernet.speed: 0
802-3-ethernet.duplex: --
802-3-ethernet.auto-negotiate: no
802-3-ethernet.mac-address: --
802-3-ethernet.cloned-mac-address: --
802-3-ethernet.generate-mac-address-mask: --
802-3-ethernet.mac-address-blacklist: --
802-3-ethernet.mtu: auto
802-3-ethernet.s390-subchannels: --
802-3-ethernet.s390-nettype: --
```

```
802-3-ethernet.s390-options:      --
802-3-ethernet.wake-on-lan:        default
802-3-ethernet.wake-on-lan-password: --
ipv4.method:                       auto
ipv4.dns:                          --
ipv4.dns-search:                   --
ipv4.dns-options:                  ""
ipv4.dns-priority:                 0
ipv4.addresses:                   --
ipv4.gateway:                     --
ipv4.routes:                      --
ipv4.route-metric:                 -1
ipv4.route-table:                  0 (unspec)
ipv4.ignore-auto-routes:           no
ipv4.ignore-auto-dns:             no
ipv4.dhcp-client-id:              --
ipv4.dhcp-timeout:                 0 (default)
ipv4.dhcp-send-hostname:           yes
ipv4.dhcp-hostname:               --
ipv4.dhcp-fqdn:                   --
ipv4.never-default:               no
ipv4.may-fail:                    yes
ipv4.dad-timeout:                  -1 (default)
ipv6.method:                      auto
ipv6.dns:                         --
ipv6.dns-search:                   --
ipv6.dns-options:                  ""
ipv6.dns-priority:                 0
ipv6.addresses:                   --
ipv6.gateway:                     --
ipv6.routes:                      --
ipv6.route-metric:                 -1
ipv6.route-table:                  0 (unspec)
ipv6.ignore-auto-routes:           no
ipv6.ignore-auto-dns:             no
ipv6.never-default:               no
ipv6.may-fail:                    yes
ipv6.ip6-privacy:                  -1 (unknown)
ipv6.addr-gen-mode:               stable-privacy
ipv6.dhcp-duid:                   --
ipv6.dhcp-send-hostname:           yes
```

```

ip6v6.dhcp-hostname:      --
ip6v6.token:              --
proxy.method:             none
proxy.browser-only:       no
proxy.pac-url:            --
proxy.pac-script:         --
GENERAL.NAME:             ens33
GENERAL.UUID:             db6f53bd-654e-45dd-97ef-
                          224514f8050a
GENERAL.DEVICES:          ens33
GENERAL.STATE:            activated
GENERAL.DEFAULT:          yes
GENERAL.DEFAULT6:         no
GENERAL.SPEC-OBJECT:      --
GENERAL.VPN:              no
GENERAL.DBUS-PATH:        /org/freedesktop/
                          NetworkManager/
                          ActiveConnection/1
GENERAL.CON-PATH:         /org/freedesktop/
                          NetworkManager/Settings/1
GENERAL.ZONE:             --
GENERAL.MASTER-PATH:      --
IP4.ADDRESS[1]:           192.168.4.210/24
IP4.GATEWAY:              192.168.4.2
IP4.ROUTE[1]:             dst = 0.0.0.0/0,
                          nh = 192.168.4.2, mt = 100
IP4.ROUTE[2]:             dst = 192.168.4.0/24,
                          nh = 0.0.0.0, mt = 100
IP4.DNS[1]:               192.168.4.2
IP4.DOMAIN[1]:            localdomain
DHCP4.OPTION[1]:          domain_name = localdomain
DHCP4.OPTION[2]:          domain_name_servers =
                          192.168.4.2
DHCP4.OPTION[3]:          expiry = 1558780229
DHCP4.OPTION[4]:          ip_address = 192.168.4.210
DHCP4.OPTION[5]:          requested_broadcast_
                          address = 1
DHCP4.OPTION[6]:          requested_dhcp_server_
                          identifier = 1
DHCP4.OPTION[7]:          requested_domain_name = 1
DHCP4.OPTION[8]:          requested_domain_name_
                          servers = 1
DHCP4.OPTION[9]:          requested_domain_search = 1

```

```

DHCP4.OPTION[10]: requested_host_name = 1
DHCP4.OPTION[11]: requested_interface_mtu = 1
DHCP4.OPTION[12]: requested_ms_classless_
                    static_routes = 1
DHCP4.OPTION[13]: requested_nis_domain = 1
DHCP4.OPTION[14]: requested_nis_servers = 1
DHCP4.OPTION[15]: requested_ntp_servers = 1
DHCP4.OPTION[16]: requested_rfc3442_classless_
                    static_routes = 1
DHCP4.OPTION[17]: requested_routers = 1
DHCP4.OPTION[18]: requested_static_routes = 1
DHCP4.OPTION[19]: requested_subnet_mask = 1
DHCP4.OPTION[20]: requested_time_offset = 1
DHCP4.OPTION[21]: requested_wpad = 1
DHCP4.OPTION[22]: routers = 192.168.4.2
DHCP4.OPTION[23]: subnet_mask = 255.255.255.0
IP6.ADDRESS[1]: fe80::959:3b1a:9607:8928/64
IP6.GATEWAY: --
IP6.ROUTE[1]: dst = fe80::/64, nh = ::,
               mt = 100
IP6.ROUTE[2]: dst = ff00::/8, nh = ::,
               mt = 256, table=255

```

To find out what exactly these settings are doing, execute **man 5 nm-settings**. You can also use **nmcli** to show an overview of currently configured devices and the status of these devices. Type, for instance, the **nmcli dev status** command to show a list of all devices, and **nmcli dev show <devicename>** to show settings for a specific device.

TIP Using **nmcli** might seem difficult. It's not, because it offers excellent command-line completion features—just make sure that the **bash-completion** package has been installed. Try it by typing **nmcli**, but don't press Enter! Instead, press the Tab key twice—you will see all available options that **nmcli** expects at this moment. Choose an option, such as **connection**, and press the Tab key twice. Using this approach helps you to compose long commands without the need to memorize anything!

In Exercise 8-3, you learn how to create connections and switch between connections using the **nmcli** command.

Exercise 8-3 Managing Network Connections with nmcli

In this exercise you create a new connection and manage its status. This connection needs to be connected to a network device. In this exercise the device `ens33` is used. If necessary, change this to the name of the network device in use on your computer.

1. Create a new network connection by typing **`nmcli con add con-name dhcp type ethernet ifname ens33 ipv4.method auto`**.
2. Create a connection with the name *static* to define a static IP address and gateway: **`nmcli con add con-name static ifname ens33 autoconnect no type ethernet ip4 10.0.0.10/24 gw4 10.0.0.1 ipv4.method manual`**. The gateway might not exist in your configuration, but that does not matter.
3. Type **`nmcli con show`** to show the connections, and use **`nmcli con up static`** to activate the static connection. Switch back to the DHCP connection using **`nmcli con up dhcp`**.

In this exercise, you created network connections using **`nmcli con add`**. You can also change current connection properties by using **`nmcli con mod`**.

In Exercise 8-4, you'll learn how to change connection parameters with **`nmcli`**.

Exercise 8-4 Changing Connection Parameters with nmcli

1. Make sure that the static connection does not connect automatically by using **`nmcli con mod static connection.autoconnect no`**.
2. Add a DNS server to the static connection by using **`nmcli con mod static ipv4.dns 10.0.0.10`**. Notice that while adding a network connection you use **`ip4`**, but while modifying parameters for an existing connection, you often use **`ipv4`** instead. This is not a typo; it is just how it works.
3. To add a second item for the same parameters, use a `+` sign. Test this by adding a second DNS server, using **`nmcli con mod static +ipv4.dns 8.8.8.8`**.
4. Using **`nmcli con mod`**, you can also change parameters such as the existing IP address. Try this by using **`nmcli con mod static ipv4.addresses 10.0.0.100/24`**.
5. And to add a second IP address, you use the `+` sign again: **`nmcli con mod static +ipv4.addresses 10.20.30.40/16`**.
6. After changing connection properties, you need to activate them. To do that, you can use **`nmcli con up static`**.

This is all you need to know about **nmcli** for the RHCSA exam. As you've noticed, **nmcli** is a very rich command. The exact syntax of this command may be hard to remember. Fortunately, though, there is an excellent man page with examples. Type **man nmcli-examples** to show this man page; you'll notice that if you can find this man page, you can do almost anything with **nmcli**. Also, don't forget to use tab completion while working with **nmcli**.

Configuring the Network with nmtui

If you do not like the complicated syntax of the **nmcli** command line, you might like **nmtui**. This is a text user interface that allows you to create network connections easily. Figure 8-2 shows what the **nmtui** interface looks like.



FIGURE 8-2 The **nmtui** Interface

The **nmtui** interface consists of three menu options:

- **Edit a Connection:** Use this option to create new connections or edit existing connections.
- **Activate a Connection:** Use this to (re)activate a connection.
- **Set System Hostname:** Use this to set the hostname of your computer.

The option to edit a connection offers almost all features that you might ever need while working on network connections. It sure allows you to do anything you need to be doing on the RHCSA exam. You can use it to add any type of connection; not just Ethernet connections, but also advanced connection types such as network bridges and teamed network drivers are supported.

When you select the option Edit a Connection, you get access to a rich interface that allows you to edit most properties of network connections. After editing the connection, you need to deactivate it and activate it again. This should work automatically, but the fact is that it does not.

TIP If you like graphical user interface (GUI) tools, you are lucky. Use `nm-connection-editor` instead of `nmtui`, but be prepared that this interface offers a relatively restricted option set. It does not contain advanced options such as the options to create network team interfaces and manage network bridge interfaces. It does, however, offer all you need to manage address configuration on a network connection. Start it by using the `nm-connection-editor` command or by using the applet in the GNOME graphical interface. Figure 8-3 shows what the default interface of this tool looks like.




FIGURE 8-3 The `nm-connection-editor` Interface

Working on Network Configuration Files

Every connection that you create is stored as a configuration file in the directory `/etc/sysconfig/network-scripts`. The name of the configuration files starts with `ifcfg-` and is followed by the name of the network interface. In Example 8-7, you can see what such a configuration file looks like.

Example 8-7 Example of an `ifcfg` Configuration File

```
[root@server1 ~]# cat /etc/sysconfig/network-scripts/ifcfg-ens33
TYPE="Ethernet"
PROXY_METHOD="none"
BROWSER_ONLY="no"
BOOTPROTO="dhcp"
DEFROUTE="yes"
IPV4_FAILURE_FATAL="no"
IPV6INIT="yes"
IPV6_AUTOCONF="yes"
IPV6_DEFROUTE="yes"
IPV6_FAILURE_FATAL="no"
IPV6_ADDR_GEN_MODE="stable-privacy"
NAME="ens33"
UUID="db6f53bd-654e-45dd-97ef-224514f8050a"
DEVICE="ens33"
ONBOOT="yes"
```

If you don't like making modifications to the network configuration with **nmcli** or **nmtui**, you can directly edit this configuration file instead. After making changes to the configuration file, use the **nmcli con up** command to activate the new configuration.

TIP You can set both a fixed IP address and a dynamic IP address in one network connection. To do that, set the `BOOTPROTO` option in the connection configuration file to **dhcp**, while you also specify an IP address and network prefix. You can do this also from the **nmtui** utility; just make sure that in **nmtui** the IPv4 configuration is set to automatic (and not to manual) and specify an IP address as well. I recommend that you do this in the test configuration you are using with this book, because it allows you to use a static network address configuration for internal use, in addition to a dynamic configuration that allows you to access the Internet and install software from repositories.

Setting Up Hostname and Name Resolution

To communicate with other hosts, hostnames are used. As an administrator, it is important that you know how to set the hostname. You also need to make sure that hosts can contact one another based on hostnames by setting up hostname resolution. In this section, you learn how to do that.

Hostnames

Because hostnames are used to access servers and the services they're offering, it is important to know how to set the system hostname. A hostname typically consists of different parts. These are the name of the host and the DNS domain in which the host resides. These two parts together make up for the fully qualified domain name (FQDN), which looks like `server1.example.com`. It is good practice to always specify an FQDN, and not just the hostname, because the FQDN provides a unique identity on the Internet. There are different ways to change the hostname:

- Use **nmtui** and select the option **Change Hostname**.
- Use **hostnamectl set-hostname**.
- Edit the contents of the configuration file `/etc/hostname`.

To configure the hostname with **hostnamectl**, you can use a command like **hostnamectl set-hostname myhost.example.com**. After setting the hostname, you can use **hostnamectl status** to show the current hostname. Example 8-8 shows the output of this command.

Example 8-8 Showing Current Hostname Configuration

```
[root@server1 ~]# hostnamectl status
Static hostname : server1.example.com
Icon name : computer-vm
Chassis : vm
Machine ID : 5aa095b495ed458d934c54a88078c165
Boot ID : 5fdef4be9cab48c59873af505d778761
Virtualization : vmware
Operating System : Red Hat Enterprise Linux 8.0 (Ootpa)
CPE OS Name : cpe:/o:redhat:enterprise_linux:8.0:GA
Kernel : Linux 4.18.0-80.el8.x86_64
Architecture : x86-64
```

When using **hostnamectl status**, you see not only information about the hostname but also information about the Linux kernel, virtualization type, and much more.

Alternatively, you can set the hostname using the **nmtui** interface. Figure 8-4 shows the screen from which this can be done.




FIGURE 8-4 Changing the Hostname Using **nmtui**

To set hostname resolution, DNS is typically used. Configuring a DNS server is not an RHCSA objective, but you need to know how to configure your server to use an existing DNS server for hostname resolution. Apart from DNS, you can configure hostname resolution in the `/etc/hosts` file. Example 8-9 shows the contents of an `/etc/hosts` file.

Example 8-9 `/etc/hosts` Sample Contents

```
[root@server1 ~]# cat /etc/hosts
127.0.0.1    localhost localhost.localdomain localhost4 localhost4.
             localhost4
::1         localhost localhost.localdomain localhost6 localhost6.
             localhost6
```

All hostname–IP address definitions as set in `/etc/hosts` will be applied before the hostname in DNS is used. This is configured as a default in the hosts line in `/etc/nsswitch.conf`, which by default looks like this:

```
hosts: files dns myhostname
```

Setting up an `/etc/hosts` file is easy; just make sure that it contains at least two columns. The first column has the IP address of the specific host, and the second column specifies the hostname. The hostname can be provided as a short name (like `server1`) or as an FQDN. In an FQDN, the hostname as well as the complete DNS name are included, as in `server1.example.com`.

If a host has more than one name, like a short name and a fully qualified DNS name, you can specify both of them in `/etc/hosts`. In that case, the second column must contain the FQDN, and the third column can contain the alias. Example 8-10 shows a hostname configuration example.

Example 8-10 /etc/hosts Configuration Example

```
[root@server2 ~]# cat /etc/hosts
127.0.0.1    localhost localhost.localdomain localhost4 localhost4.
             localhost4
::1         localhost localhost.localdomain localhost6 localhost6.
             localhost6
10.0.0.10    server1.example.com server1
10.0.0.20    server2.example.com server2
```

DNS Name Resolution

Just using an /etc/hosts file is not enough for name resolution if you want to be able to communicate with other hosts on the Internet. You should use DNS, too. To specify which DNS server should be used, set the DNS server using **nmcli** or **nmtui** as previously discussed. The NetworkManager configuration stores the DNS information in the configuration file for the network connection, which is in /etc/sysconfig/network-scripts, and from there pushes the configuration to the /etc/resolv.conf file, which is used for DNS name server resolving. Do not edit /etc/resolv.conf directly, as it will be overwritten the next time you restart NetworkManager.

It is recommended to always set up at least two DNS name servers to be contacted. If the first name server does not answer, the second name server is contacted. To specify which DNS name servers you want to use, you have a few different options:

- Use **nmtui** to set the DNS name servers. Figure 8-5 shows the interface from which you can do this.
- Set the DNS1 and DNS2 parameters in the ifcfg network connection configuration file in /etc/sysconfig/network-scripts.
- Use a DHCP server that is configured to hand out the address of the DNS name server.
- Use **nmcli con mod <connection-id> [+]**ipv4.dns** <ip-of-dns>**.




FIGURE 8-5 Setting DNS Servers from the **nmtui** Interface

Notice that if your computer is configured to get the network configuration from a DHCP server, the DNS server is also set via the DHCP server. If you do not want this to happen, you have two options:

- Edit the `ifcfg` configuration file to include the option **PEERDNS=no**.
- Use **nmcli con mod <con-name> ipv4.ignore-auto-dns yes**.

To verify hostname resolution, you can use the **getent hosts <servername>** command. This command searches in both `/etc/hosts` and DNS to resolve the hostname that has been specified.

EXAM TIP Do *not* specify the DNS servers directly in `/etc/resolv.conf`. They will be overwritten by NetworkManager when it is (re)started.

Summary

In this chapter, you learned how to configure networking in RHEL 8. First you read how the IP protocol is used to connect computers, and then you read which techniques are used to make services between hosts accessible. Next you read how to verify the network configuration using the **ip** utility and some related utilities. In the last part of this chapter, you read how to set IP addresses and other host configuration in a permanent way by using either the **nmcli** or the **nmcli** utility.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 8-3 lists a reference of these key topics and the page number on which each is found.



Table 8-3 Key Topics for Chapter 8

Key Topic Element	Description	Page
List	IPv4 and IPv6 short descriptions	170
List	Private network addresses	170
Table 8-2	Binary-decimal conversion overview	172
List	IP address types	174

Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

IP, IPv4, IPv6, protocol, port, subnet mask, DNS, DHCP, connection, interface, FQDN

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. What is the network address in the address 213.214.215.99/29?
2. Which command only shows link status and not the IP address?
3. Which service manages network configuration in RHEL 8?
4. Which file contains the hostname in RHEL 8?
5. Which command enables you to set the hostname in an easy way?
6. Which command do you need to run after manually changing the contents of the `/etc/sysconfig/ifcfg` files?
7. Which configuration file can you change to enable hostname resolution for a specific IP address?
8. Which command shows the current routing configuration?
9. How do you verify the current status of the NetworkManager service?
10. Which command enables you to change the current IP address and default gateway on your network connection?

End-of-Chapter Lab

For exercises in later chapters in this book, it is recommended to have a test environment in which at least two servers are present. To do the exercises in this lab, make sure that you have a second server installed.

Lab 8.1

1. Set up the first server to use the FQDN **server1.example.com**. Set up the second server to use **server2.example.com**.
2. On **server1.example.com**, use **nmcli** and configure your primary network card to automatically get an IP address through DHCP. Also set a fixed IP address to **192.168.4.210**. On **server2**, set the fixed IP address to **192.168.4.220**.
3. Make sure that from **server1** you can ping **server2**, and vice versa.
4. To allow you to access servers on the Internet, make sure that your local DHCP server provides the default router and DNS servers.



The following topics are covered in this chapter:

- Managing Software Packages with YUM
- Using **yum**
- Managing Package Module Streams
- Managing Software Packages with RPM

The following RHCSA exam objectives are covered in this chapter:

- Install and update software packages from Red Hat Network, a remote repository, or from the local file system
- Work with package module streams

Managing software packages is an important task for an administrator of Red Hat Enterprise Linux. In this chapter, you learn how to manage software packages from the command line by using the **yum** utility. You also learn which role repositories play in software management with YUM. Next we'll cover working with Package Module Streams, a solution that makes it possible to work with the specific version packages that you need in your environment. In the last part of this chapter, you learn how to manage software with the **rpm** command, which is useful to query new and installed software packages.

Managing Software

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 9-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 9-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Managing Software Packages with yum	1–4
Using yum	5
Managing Package Module Streams	6–7
Managing Software Packages with rpm	8–10

- Which of the following is *not* a mandatory component in a .repo file that is used to indicate which repositories should be used?
 - [label]
 - name=
 - baseurl=
 - gpgcheck=
- Which installation source is used on RHEL if a server is not registered with RHN?
 - The installation medium is used.
 - No installation source is used.
 - The base RHN repository is used, without updates.
 - You have full access to RHN repositories, but the software you are using is not supported.

3. Which of the following should be used in the `.repo` file to refer to a repository that is in the directory `/repo` on the local file system?
 - a. `file=/repo`
 - b. `baseurl=file://repo`
 - c. `baseurl=file:///repo`
 - d. `file=http:///repo`
4. Which of the following is true about GPG-based repository security?
 - a. If packages in the repository have been signed, you need to copy the GPG key to the correct location.
 - b. GPG package signing is mandatory.
 - c. GPG package signatures prevent packages in a repository from being changed.
 - d. GPG package signing is recommended on Internet repositories, but not required on local repositories that are for internal use only.
5. Which command enables you to search the package that contains the file `semanage`?
 - a. **`yum search semanage`**
 - b. `yum search all semanage`
 - c. `yum provides semanage`
 - d. `yum whatprovides */semanage`
6. Which YUM module component allows you to work with different versions side by side?
 - a. Application profile
 - b. Application stream
 - c. Module version
 - d. RPM group
7. Which of the following commands allows you to install the devel profile of the PHP 7.1 application stream?
 - a. **`yum module install php:7.1 devel`**
 - b. `yum module install php:7.1 --devel`
 - c. `yum module install php:7.1/devel`
 - d. `yum module install php:7.1@devel`

8. Which command should you use to install an RPM file that has been downloaded to your computer?
 - a. `yum install`
 - b. `yum localinstall`
 - c. `rpm -ivh`
 - d. `rpm -Uvh`
9. Which command enables you to find the RPM package a specific file belongs to?
 - a. `rpm -ql /my/file`
 - b. `rpm -qlf /my/file`
 - c. `rpm -qf /my/file`
 - d. `rpm -qa /my/file`
10. Which command enables you to analyze whether there are scripts in an RPM package file that you have just downloaded?
 - a. `rpm -qs packagename.rpm`
 - b. `rpm -qps packagename.rpm`
 - c. `rpm -qp --scripts packagename.rpm`
 - d. `rpm -q --scripts packagename.rpm`

Foundation Topics

Managing Software Packages with yum

The default utility used to manage software packages on Red Hat Enterprise Linux is `yum`, which stands for Yellowdog Updater, Modified. `yum` is designed to work with repositories, which are online depots of available software packages. In this section, you learn how to create and manage repositories and how to manage software packages based on the contents of the repositories.

NOTE Software in RHEL is based on Fedora Software. In the versions of Fedora that RHEL 8 is based on, `yum` recently has been replaced with the `dnf` utility. For that reason, it was expected that with the release of RHEL 8, the `yum` command would be replaced with the `dnf` command. Red Hat has decided differently. With RHEL 8, a new version of `yum` has been introduced, which is based on the `dnf` command. You'll notice that in many cases, when requesting information about `yum`, you're redirected to `dnf` resources. So in fact you're using `dnf`, but Red Hat has decided to rename it to `yum`. Even if in reality you're using the `dnf` command, in this book I'm following the Red Hat convention, and for that reason you'll learn how to work with `yum`.

Understanding the Role of Repositories

Software on Red Hat Enterprise Linux is provided in the RPM (Red Hat Package Manager) format. This is a specific format used to archive the package and provide package metadata as well.

When you are working with software in RHEL, repositories play a key role. Working with repositories makes it easy to keep your server current: The maintainer of the repository publishes updated packages in the repository, and the result is that whenever you use the `yum` command (discussed later in this chapter) to install software, the most recent version of the software is automatically used.

Another major benefit of working with YUM is the way that package dependencies are dealt with. On Linux (as on most other modern operating systems), software packages have dependencies. This means that to install one package, other packages have to be present as well. Without using repositories, that would mean that these packages have to be installed manually.

The YUM repository system takes care of resolving these dependencies automatically. If a package is going to be installed, it contains information about the required dependencies. The `yum` command then looks in the repositories configured on this system to fetch the dependencies automatically. If all goes well, the installer just sees

a short list of the dependencies that will be installed as a dependency to install the package. If you are using RHEL with the repositories that are provided through Red Hat Network (RHN), there is no reason why this procedure should not work, and the attempts to install software will usually succeed.

While installing RHEL 8, it asks you to register with the Red Hat Customer Portal, which provides different repositories. After registering, you can install software packages that are verified by Red Hat automatically. If you choose to install RHEL without registration, it cannot get in touch with the Red Hat repositories, and you end up with no repositories at all. In that case, you have to be able to specify yourself which repository you want to use. If you are using CentOS, you get access to the CentOS repositories and no registration is required.

Note that repositories are specific to an operating system. Therefore, if you are using RHEL, you should use RHEL repositories only. Do not try, for instance, to add CentOS repositories to a RHEL server. If you want to provide additional software from the Fedora project to a RHEL server (which for support reasons is not recommended), you can consider adding the EPEL (Extra Packages for Enterprise Linux) repositories. See <https://fedoraproject.org/wiki/EPEL> for more information.

WARNING Before adding the EPEL repository to RHEL, make sure that it doesn't break your current support status. EPEL packages are not managed by Red Hat, and adding them make break supported Red Hat packages.

Registering Red Hat Enterprise Linux for Support

Red Hat Enterprise Linux is a supported Linux operating system that requires you to register. To register RHEL, you need a valid entitlement. This entitlement is associated to your account on the Red Hat Customer Portal. You can obtain an entitlement by purchasing a subscription for RHEL or by joining the Red Hat Developer program, which gives access to the no-cost Red Hat Enterprise Developer subscription. You can sign up for the Red Hat Developer subscription at <https://developers.redhat.com>.

After obtaining a valid subscription for Red Hat Enterprise Linux, you can use the Red Hat Subscription Management (RHSM) tools to manage your entitlement. Managing an entitlement involves four basic tasks:

- **Register:** While registering a subscription, you connect it to your current Red Hat account. As a result, the **subscription-manager** tool can inventory the system. If a system is no longer used, it can also be unregistered.

- **Subscribe:** Subscribing a system gives it access to updates for Red Hat products that your subscription is entitled to. Also, by subscribing, you'll get access to the support level that is associated with your account.
- **Enable repositories:** After subscribing a system, you'll get access to a default set of repositories. Some repositories by default are disabled but can be enabled after subscribing your system.
- **Review and track:** You can review and track current subscriptions that are in use.

Managing Subscriptions

You can manage subscriptions either from the GNOME graphical interface or from the command line. The **subscription-manager** tool is used for managing subscriptions from the command line. You can use it in the following ways:

- **Register a system:** Type **subscription-manager register** to register. It will prompt for the name of your Red Hat user account as well as your password, and after entering these, your RHEL server will be registered.
- **List available subscriptions:** Each account has access to specific subscriptions. Type **subscription-manager list --available** to see what your account is entitled to.
- **Automatically attach a subscription:** Registering a server is not enough to get access to the repositories. Use **subscription-manager attach --auto** to automatically attach your subscription to the repositories that are available.
- **Get an overview:** To see which subscriptions you're currently using, type **subscription-manager --consumed**.
- **Unregister:** If you're going to deprovision a system, use **subscription-manager unregister**. If you have access to a limited number of registered systems only, unregistering is important to ensure that you don't run out of available licenses.

After registering and attaching a subscription, entitlement certificates are written to the `/etc/pki` directory. In `/etc/pki/product`, certificates are stored that indicate which Red Hat products are installed on this system. In `/etc/pki/consumer`, certificates are stored that identify the Red Hat account to which the system is registered, and the `/etc/pki/entitlement` directory contains information about the subscriptions that are attached to this system.

Specifying Which Repository to Use

On most occasions, after the installation of your server has finished, it is configured with a list of repositories that should be used. You sometimes have to tell your server which repositories should be used:

- You want to distribute nondefault software packages through repositories.
- You are installing Red Hat Enterprise Linux without registering it on RHN.

Telling your server which repository to use is not difficult, but it is important that you know how to do it (for the RHCSA exam, too).

IMPORTANT! To learn how to work with repositories and software packages, do *not* use the repositories that are provided by default. So if you have installed RHEL, do *not* register using **subscription-manager**, and if you have installed CentOS, remove all files from `/etc/yum.repos.d`.

**Key
Topic**

To tell your server which repository to use, you need to create a file with a name that ends in `.repo` in the directory `/etc/yum.repos.d`. In that file you need the following contents:

- **[label]** The `.repo` file can contain different repositories, each section starting with a label that identifies the specific repository.
- **name=** Use this to specify the name of the repository you want to use.
- **baseurl=** Contains the URL that points to the specific repository location.

In the repository files that are provided by default, you may find several repositories in one file, as is the case in Example 9-1. This is useful to group repositories that belong together in one file, and is often done in repository files that are provided as a default. If you are creating repository files yourself, you are free to create separate files for each repository.

Example 9-1 shows a repository file that is based on the default repositories that are installed on CentOS 8.

Example 9-1 Repository File Example

```
[root@server1 yum.repos.d]# cat CentOS-Base.repo
# CentOS-Base.repo
#
# The mirror system uses the connecting IP address of the client and
# the
# update status of each mirror to pick mirrors that are updated to
# and
# geographically close to the client. You should use this for CentOS
# updates
# unless you are manually picking other mirrors.
#
# If the mirrorlist= does not work for you, as a fall back you can
# try the
# remarked out baseurl= line instead.
#
#

[base]
name=CentOS-$releasever - Base
mirrorlist=http://mirrorlist.centos.org/?release=$releasever&arch=
    $basearch&repo=os
#baseurl=http://mirror.centos.org/centos/$releasever/os/$basearch/
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-8

#released updates
[updates]
name=CentOS-$releasever - Updates
mirrorlist=http://mirrorlist.centos.org/?release=$releasever&arch=
    $basearch&repo=updates
#baseurl=http://mirror.centos.org/centos/$releasever/updates/$basearch/
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-8

#additional packages that may be useful
[extras]
name=CentOS-$releasever - Extras
mirrorlist=http://mirrorlist.centos.org/?release=$releasever&arch=
    $basearch&repo=extras
#baseurl=http://mirror.centos.org/centos/$releasever/extras/$basearch/
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-8
```

```
#additional packages that extend functionality of existing packages
[centosplus]
name=CentOS-$releasever - Plus
mirrorlist=http://mirrorlist.centos.org/?release=$releasever&arch=
    $basearch&repo=centosplus
#baseurl=http://mirror.centos.org/centos/$releasever/centosplus/
    $basearch/
gpgcheck=1
enabled=0
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-8
```

In the repository configuration file from Example 9-1, you can see that some options are used. Table 9-2 summarizes these options.

Table 9-2 Key Options in .repo Files

Option	Explanation
[label]	The label used as an identifier in the repository file.
name=	Specifies the name of the repository.
mirrorlist=	Refers to a URL where information about mirror servers for this server can be obtained. Typically used for big online repositories only.
baseurl=	Refers to the base URL where the RPM packages are found.
gpgcheck=	Set to 1 if a GNU Privacy Guard (GPG) integrity check needs to be performed on the packages. If set to 1, a GPG key is required.
gpgkey=	Specifies the location of the GPG key that is used to check package integrity.

When creating a repository file, the `baseurl` parameter is the most important because it tells your server where to find the files that are to be installed. The `baseurl` takes as its argument the URL where files need to be installed from. This will often be an HTTP or FTP URL, but it can be a file-based URL as well.

When using a URL, two components are used. First, the URL identifies the protocol to be used and is in the format `protocol://`, such as `http://`, `ftp://`, or `file://`. Following the URL is the exact location on that URL. That can be the name of a web server or an FTP server, including the subdirectory where the files are found. If the URL is file based, the location on the file system starts with a `/` as well.

Therefore, for a file system-based URL, there will be three slashes in the `baseurl`, such as `baseurl:///repo`, which refers to the directory `/repo` on the local file system.

Understanding Repository Security

Using repositories allows you to transparently install software packages from the Internet. This is convenient, but it also involves a security risk. When installing RPM packages, you do that with root permissions, and if in the RPM package script code is executed (which is common), that is executed as root as well. For that reason, you want to make sure that you can trust the software packages you are trying to install. This is why repositories in general use keys for package signing. This is also why on Red Hat Enterprise Linux it is a good idea to use repositories provided through RHN only.

To secure packages in a repository, these packages are often signed with a GPG key. This makes it possible to check whether packages have been changed since the owner of the repository provided them. The GPG key used to sign the software packages is typically made available through the repository as well. The users of the repository can download that key and store it locally so that the package signature check can be performed automatically each time a package is downloaded from the repository.

If repository security is compromised and an intruder manages to hack the repository server and put some fake packages on it, the GPG key signature will not match, and the **yum** command will complain while installing new packages. This is why it is highly recommended to use GPG keys when using Internet repositories.

If you are using a repository where GPG package signing has been used, on first contact with that repository, the **rpm** command will propose to download the key that was used for package signing (see Example 9-2). This is a transparent procedure that requires no further action. The GPG keys that were used for package signing are installed to the `/etc/pki/rpm-gpg` directory by default.

TIP For using internal repositories, the security risks are not that high. For that reason, you do not have to know how to work with GPG-signed packages on the exam.

Example 9-2 On First Contact with a Repository, the GPG Key Is Downloaded

```
[root@server1 ~]# yum install kernel
Loaded plugins: fastestmirror
Loading mirror speeds from cached hostfile
* base: centos.mirror1.spango.com
* extras: mirror.netrouting.net
* updates: mirrors.supportex.net
Resolving Dependencies
--> Running transaction check
---> Package kernel.x86_64 0:4.10.0-229.1.2.el8 will be installed
--> Processing Dependency: linux-firmware >= 20190911 for package:
    kernel-4.10.0-229.1.2.el7.x86_64
```

```
--> Running transaction check
---> Package linux-firmware.noarch 0:20140213-0.3.git4164c23.el8
    will be updated
---> Package linux-firmware.noarch 0:20140911-0.1.git365e80c.el8
    will be an update
--> Finished Dependency Resolution

Dependencies Resolved

=====
=====
Package                               Arch                               Version
Repository                            Size
=====
Installing:
    kernel                               x86_64                               updates
3.10.0-229.1.2.el8
31M
Updating for dependencies:
    linux-firmware                       noarch
20190911-0.1.git365e80c.el8             base
17 M

Transaction Summary
=====
Install 1 Package
Upgrade      ( 1 Dependent package)

Total size: 48 M
Is this ok [y/d/N]: y
Downloading packages:
warning: /var/cache/yum/x86_64/8/base/packages/linux-firmware-20140911-
0.1.git365e80c.el7.noarch.rpm: Header V3 RSA/SHA256 Signature, key
ID f4a80eb5: NOKEY
Retrieving key from file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-8
Importing GPG key 0xF4A80EB5:
  Userid : "CentOS-8 Key (CentOS 7 Official Signing Key)
    <security@centos.org>"
  Fingerprint: 6341 ab27 53d7 8a78 a7c2 7bb1 24c6 a8a7 f4a8 0eb5
  Package : centos-release-8-0.1406.el7.centos.2.3.x86_64 (@anaconda)
  From : /etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-8
Is this ok [y/N]: y
```

Creating Your Own Repository

Creating your own repository is not a requirement for the RHCSA exam, but knowing how to do so is useful if you want to test setting up and working with repositories. Also, while doing the exercises in this book, it is important to know how to set up your own repositories, as you should not use the repositories that are available online.

Setting up your own repository allows you to put your own RPM packages in a directory and publish that directory as a repository. It is also useful to know how to do this if you have installed RHEL and not connected it to RHN, which means that you would not have any repositories at all.

The procedure itself is not hard to summarize. You need to make sure all RPM packages are available in the directory that you want to use as a repository, and after doing that, you need to use the **createrepo** command to generate the metadata that enables you to use that directory as a repository. Exercise 9-1 describes how to create your own repository.

Exercise 9-1 Creating Your Own Repository

To perform this exercise, you need to have access to the RHEL or CentOS installation disk or ISO file.

1. Insert the installation disk in your virtual machine and make sure it is attached and available.
2. Type **mkdir /repo** so that you have a mount point where you can mount the ISO file.
3. Add the following line to the end of the `/etc/fstab` configuration file: **/dev/sr0 /repo iso9660 defaults 0 0**
4. Type **mount -a**, followed by **mount | grep sr0**. You should now see that the optical device is mounted on the directory `/repo`. At this point, the directory `/repo` can be used as a repository.
5. At this point, two repositories are available through the `/repo` directory. The BaseOS repository provides access to the base packages, and the Application Stream (AppStream) repository provides access to application streams (these repositories are described in more detail in the “Managing Package Module Streams” section later in the chapter). To make them accessible, you need to add two files to the `/etc/yum.repos.d` directory. Start with the file `BaseOS` and give it the following contents:

```
[BaseOS]
name=BaseOS
baseurl=file:///repo/BaseOS
gpgcheck=0
```

6. Next, add the file `/etc/yum.repos.d/AppStream.repo` with the following contents:


```
[AppStream]
name=AppStream
baseurl=file:///repo/AppStream
gpgcheck=0
```
7. Type **yum repolist** to verify the availability of the newly created repository. It should show the name of the myrepo repository, including the number of packages offered through this repository (see Example 9-3). Notice that if you're doing this on RHEL, you'll also see a message that this system is not registered to Red Hat Subscription Management. You can safely ignore that message.

Example 9-3 Verifying Repository Availability with **yum repolist**

```
[root@server3 yum.repos.d]# yum repolist
Updating Subscription Management repositories.
Unable to read consumer identity
This system is not registered to Red Hat Subscription Management.
You can use subscription-manager to register.
AppStream                               194 MB/s | 5.3 MB      00:00
BaseOS                                  86 MB/s | 2.2 MB      00:00
Last metadata expiration check: 0:00:01 ago on Sat 25 May 2019
09:20:19 AM EDT.
repo id                                repo name                                status
AppStream                              AppStream                                4,672
BaseOS                                  BaseOS                                    1,658
```

Using yum

At this point, you should have operational repositories, so it is time to start using them. To use repositories, you need the **yum** command. This command enables you to perform several tasks on the repositories. Table 9-3 provides an overview of common **yum** tasks.

Key
Topic

Table 9-3 Common **yum** Tasks

Task	Explanation
search	Search for the exact name of a package
[what]provides */name	Perform a deep search in the package to look for specific files within the package
info	Provide more information about the package
install	Install the package
remove	Remove the package
list [all installed]	List all or installed packages
group list	List package groups
group install	Install all packages from a group
update	Update packages specified
clean all	Remove all stored metadata

Using **yum** to Find Software Packages

To install packages with **yum**, you first need to know the exact name of the package. The **yum search** command can help you with that. When you use **yum search**, it first gets in touch with the online repositories (which might take a minute), after which it downloads the most recent repository metadata to the local machine. Then, **yum search** looks in the package name and description for the string you have been looking for. In Example 9-4, you can see what the result looks like after using **yum search user**.

Example 9-4 **yum search** Sample Output

```
[root@server3 ~]# yum search user
Updating Subscription Management repositories.
Unable to read consumer identity
This system is not registered to Red Hat Subscription Management.
You can use subscription-manager to register.
Last metadata expiration check: 0:02:23 ago on Sat 25 May 2019
09:20:19 AM EDT.
===== Name & Summary Matched: user =====
trousers-lib.x86_64      : TrouSerS libtspi library
trousers-lib.i686       : TrouSerS libtspi library
trousers-lib.x86_64     : TrouSerS libtspi library
gnome-user-docs.noarch  : GNOME User Documentation
gnome-user-docs.noarch  : GNOME User Documentation
xdg-user-dirs.x86_64    : Handles user special directories
xdg-user-dirs.x86_64    : Handles user special directories
util-linux-user.x86_64  : libuser based util-linux utilities
...
```

Because the **yum search** command looks in the package name and summary only, it often does not show what you need. You often need to look for packages containing a specific file. To do this, the **yum whatprovides** command or **yum provides** command will help you. (There is no functional difference between these two commands.) To make it clear that you are looking for packages containing a specific file, you need to specify the filename as ***/filename**, or use the full path name to the file you want to use. So if you need to look for the package containing the file **semanage**, for example, use **yum whatprovides */semanage**. It will show the name of the package as a result. Try it, and compare the result to the result of the command **yum search semanage**; you'll notice that it is quite different.

Getting More Information About Packages

Before installing a package, it is a good idea to get some more information about the package. Because the **yum** command was developed to be intuitive, it is almost possible to guess how that works. Just use **yum info**, followed by the name of the package. In Example 9-5, you see what this looks like for the **nmap** package (which, by the way, is a very useful tool). It is a network sniffer that allows you to find ports that are open on other hosts. Just use **nmap 192.168.4.100** to give it a try, but be aware that some network administrators really do not like **nmap** and might consider this a hostile attack.

Example 9-5 Example Output of **yum info nmap**

```
[root@server3 ~]# yum info nmap
Updating Subscription Management repositories.
Unable to read consumer identity
This system is not registered to Red Hat Subscription Management.
You can use subscription-manager to register.
Last metadata expiration check: 0:05:01 ago on Sat 25 May 2019
09:20:19 AM EDT.
Available Packages
Name           : nmap
Epoch         : 2
Version        : 7.70
Release        : 4.el8
Arch           : x86_64
Size           : 5.8 M
Source          : nmap-7.70-4.el8.src.rpm
Repo           : AppStream
Summary        : Network exploration tool and security scanner
URL            : http://nmap.org/
License        : Nmap
```

```

Description      : Nmap is a utility for network exploration or security
                  auditing.  It supports
                  : ping scanning (determine which hosts are up), many
                  : port scanning techniques
                  : (determine what services the hosts are offering), and
                  : TCP/IP fingerprinting
                  : (remote host operating system identification). Nmap
                  : also offers flexible target
                  : and port specification, decoy scanning, determination
                  : of TCP sequence
                  : predictability characteristics, reverse-identd
                  : scanning, and more. In addition
                  : to the classic command-line nmap executable, the Nmap
                  : suite includes a flexible
                  : data transfer, redirection, and debugging tool (netcat
                  : utility ncat), a utility
                  : for comparing scan results (ndiff), and a packet
                  : generation and response
                  : analysis tool (nping).

```

Installing and Removing Software Packages

If after looking at the **yum info** output you are happy with the package, the next step is to install it. As anything else you are doing with **yum**, it is not hard to guess how to do that: use **yum install nmap**. When used in this way, the **yum** command asks for confirmation. If when you type the **yum install** command you are sure about what you are doing, you might as well use the **-y** option, which passes a “yes” to the confirmation prompt that **yum** normally issues. Example 9-6 shows what the result looks like.

Example 9-6 Installing Software with **yum**

```

[root@server3 ~]# yum install nmap
Updating Subscription Management repositories.
Unable to read consumer identity
This system is not registered to Red Hat Subscription Management.
You can use subscription-manager to register.
Last metadata expiration check: 0:06:16 ago on Sat 25 May 2019
09:20:19 AM EDT.
Dependencies resolved.

```

```

=====
Package      Arch          Version      Repository    Size
=====
Installing:
  nmap        x86_64        2:7.70-4.el8  AppStream     5.8 M

Transaction Summary
=====
Install 1 Package

Total size: 5.8 M
Installed size: 24 M
Is this ok [y/N]: y
Downloading Packages:
Running transaction check
Transaction check succeeded.
Running transaction test
Transaction test succeeded.
Running transaction
  Preparing           :                               1/1
  Installing          : nmap-2:7.70-4.el8.x86_64      1/1
  Running scriptlet   : nmap-2:7.70-4.el8.x86_64      1/1
  Verifying           : nmap-2:7.70-4.el8.x86_64      1/1
Installed products updated.

Installed:
  nmap-2 :7.70-4.el8.x86_64

Complete!

```

In Example 9-6, you can see that **yum** starts by analyzing what is going to be installed. Once that is clear, it gives an overview of the package that is going to be installed, including its dependencies. Then, the package itself is installed to the system.

To remove software packages from a machine, use the **yum remove** command. This command also does a dependency analysis, which means that it will remove not only the selected package but also all packages that depend on it. This may sometimes lead to a long list of software packages that are going to be removed. To avoid unpleasant surprises, you should never use **yum remove** with the **-y** option.

NOTE Some packages are protected. Therefore, you cannot easily remove them. If **yum remove** encounters protected packages, it refuses to remove them.

Showing Lists of Packages

When working with **yum**, you may also use the **yum list** command to show lists of packages. Used without arguments, **yum list** shows a list of all software packages that are available, including the repository they were installed from. You see the repository names as listed in Table 9-2 and the **@anaconda** repository as well. If a repository name is shown, the package is available in that specific repository. If **@anaconda** is listed, the package has already been installed on this system. Example 9-7 shows the partial output of the **yum list** command.

Example 9-7 Partial Output of the **yum list** Command

```
[root@server3 ~]# yum list | less

Updating Subscription Management repositories.
Unable to read consumer identity
This system is not registered to Red Hat Subscription Management.
  You can use subscription-manager to register.
Last metadata expiration check: 0:09:31 ago on Sat 25 May 2019
 09:20:19 AM EDT.

Installed Packages
GConf2.x86_64                               3.2.6-22.el8           @AppStream
ModemManager.x86_64                         1.8.0-1.el8            @anaconda
ModemManager-glib.x86_64                   1.8.0-1.el8            @anaconda
NetworkManager.x86_64                      1:1.14.0-14.el8        @anaconda
NetworkManager-adsl.x86_64                 1:1.14.0-14.el8        @anaconda
NetworkManager-bluetooth.x86_64           1:1.14.0-14.el8        @anaconda
NetworkManager-config-server.noarch        1:1.14.0-14.el8        @anaconda
NetworkManager-libnm.x86_64                1:1.14.0-14.el8        @anaconda
NetworkManager-team.x86_64                 1:1.14.0-14.el8        @anaconda
NetworkManager-tui.x86_64                  1:1.14.0-14.el8
...
```

If you want to see which packages are installed on your server, you can use the **yum list installed** command. The **yum list** command can also prove useful when used with the name of a specific package as its argument. For instance, type **yum list kernel** to show which version of the kernel is actually installed and which version is available as the most recent version in the repositories. Example 9-8 shows the result of this command.

Example 9-8 Use **yum list *packagename*** for Information About Installed and Available Versions

```
[root@server3 ~]# yum list kernel
Updating Subscription Management repositories.
Unable to read consumer identity
This system is not registered to Red Hat Subscription Management.
You can use subscription-manager to register.
Last metadata expiration check: 0:11:57 ago on Sat 25 May 2019
09:20:19 AM EDT.
Installed Packages
kernel.x86_64                               4.18.0-80.el8                                @anacond
```

Updating Packages

One of the major benefits of working with **yum** repositories is that repositories make it easy to update packages. The individual who maintains a repository is responsible for copying updated packages to the repositories. The index in the repository always contains the current version of a package in the repository. On the local machine also, a database is available with the current versions of the packages that are used. When using the **yum update** command, current versions of packages that are installed are compared to the version of these packages in the repositories. As shown in Example 9-9, **yum** next shows an overview of updatable packages. From this overview, type **y** to install the updates.

Notice that while updating packages the old version of the package is replaced with a newer version of the package. There is one exception, which is for the kernel package. Even if you are using the **yum update kernel** command, the kernel package is not updated, but the newer kernel is installed beside the old kernel, so that while booting you can select the kernel that you want to use. This is useful if the new kernel won't work because of hardware compatibility issues. In that case, you can interrupt the GRUB 2 boot process (see Chapter 17, “Managing and Understanding the Boot Procedure,” for more details) to start the older kernel.

Example 9-9 Using **yum update**

```
[root@server1 ~]# yum update kernel
Updating Subscription Management repositories.
Red Hat Enterprise Linux 8 for x86_64 - AppStream (RPMs)
 2.4 kB/s | 4.5 kB      00:01
Red Hat Enterprise Linux 8 for x86_64 - BaseOS (RPMs)
 2.5 kB/s | 4.0 kB      00:01
```

```

created by dnf config-manager from file:///repo
0.0 B/s | 0 B 00:00
Failed to synchronize cache for repo 'repo', ignoring this repo.
Dependencies resolved.
=====
Package           Arch      Version              Repository            Size
=====
Installing:
kernel            x86_64    4.18.0-80.1.2.el8_0  rhel-8-for-x86_64-
                    baseos-rpms              406 k
Installing dependencies:
kernel-modules    x86_64    4.18.0-80.1.2.el8_0  rhel-8-for-x86_64-
                    baseos-rpms              20 M
kernel-core       x86_64    4.18.0-80.1.2.el8_0  rhel-8-for-x86_64-
                    baseos-rpms              24 M

Transaction Summary
=====
Install 3 Packages

Total download size: 44 M
Installed size: 76 M
Is this ok [y/N]: y

```

Working with yum Package Groups

While managing specific services on a Linux machine, you often need several different packages. If, for instance, you want to make your machine a virtualization host, you need the KVM packages, but also all supporting packages such as `qemu`, `libvirt`, and the client packages. Or while configuring your server as a web server, you need to install additional packages like PHP as well in many cases.

To make it easier to manage specific functionality, instead of specific packages, you can work with package groups as well. A package group is defined in the repository, and `yum` offers the group management commands to work with these groups. For an overview of all current groups, use **`yum groups list`**. This shows output as in Example 9-10.

TIP The name of the command is **`yum groups`**, but there are aliases that ensure that **`yum group`** and even commands like **`yum groupinstall`** are also working. So, you can use either of these commands.

Example 9-10 Showing Available **yum** Groups

```
[root@server1 ~]# yum groups list
Updating Subscription Management repositories.
created by dnf config-manager from file:///repo
0.0 B/s | 0 B 00:00
Failed to synchronize cache for repo 'repo', ignoring this repo.
Last metadata expiration check: 0:01:08 ago on Sat 25 May 2019
09:34:12 AM EDT.
Available Environment Groups:
  Server
  Minimal Install
  Workstation
  Virtualization Host
  Custom Operating System
Installed Environment Groups:
  Server with GUI
Installed Groups:
  Container Management
  Headless Management
Available Groups:
  RPM Development Tools
  Smart Card Support
  .NET Core Development
  Network Servers
  Development Tools
  System Tools
  Graphical Administration Tools
  Scientific Support
  Security Tools
  Legacy UNIX Compatibility
```

Notice that some yum groups are not listed by default. To show those as well, type **yum groups list hidden**. You see that the list of groups that is displayed is considerably longer. The difference is that **yum groups list** shows environment groups, which contain basic functionality. Within an environment group, different subgroups can be used; these are displayed only when using **yum groups list hidden**.

To get information about packages available in a group, you use **yum groups info**. Because group names normally contain spaces, do not forget to put the entire group name between quotes. So, type **yum groups info "Basic Web Server"** to see what is in the Basic Web Server group. As you can see in Example 9-11, this command shows mandatory items and optional items in the group. The items can be groups and individual packages.

Example 9-11 Showing Group Contents with **yum groups info**

```
[root@server1 ~]# yum groups info "Basic Web Server"
Updating Subscription Management repositories.
created by dnf config-manager from file:///repo
0.0 B/s | 0 B 00:00
Last metadata expiration check: 0:02:55 ago on Sat 25 May 2019
09:34:12 AM EDT.

Group: Basic Web Server
Description: These tools allow you to run a Web server on the
system.
Mandatory Packages:
    httpd
Default Packages:
    httpd-manual
    mod_fcgid
    mod_ssl
Optional Packages:
    libmemcached
    memcached
    mod_auth_gssapi
    mod_security
    mod_security-mlogc
    mod_security_crs
```

Using yum History

While working with YUM, all actions are registered. You can use the **yum history** command to get an overview of all actions that have been issued. From the history file, it is possible to undo specific actions; use **yum history undo** followed by the number of the specific action you want to undo.

In Example 9-12, you see the result of the **yum history** command, where every action has its own ID.

Example 9-12 Showing Past yum Actions Using **yum history**

```
[root@server1 ~]# yum history
Updating Subscription Management repositories.
ID | Command line | Date and time | Action(s) | Altered
-----
3 | install dnf-utils | 2019-05-25 09:07 | Install | 1
2 | install stratis-cli stra | 2019-05-24 07:13 | Install | 9
1 | | 2019-05-19 06:20 | Install | 1336 EE
```

As you can see, action number 2 altered nine packages and was used to install packages. To undo this action completely, type **yum history undo 2**. In Exercise 9-2, you apply some of the most useful **yum** commands for common package management tasks, as discussed previously.

Exercise 9-2 Using yum for Package Management

1. Type **yum repolist** to show a list of the current repositories that your system is using.
2. Type **yum search seinfo**. This will give no matching result.
3. Type **yum provides */seinfo**. The command shows that the `setools-console-<version>` package contains this file.
4. Install this package using **yum install -y setools-console**. Depending on your current configuration, you might notice that quite a few dependencies have to be installed also.
5. Type **yum list setools-console**. You see that the package is listed as installed.
6. Type **yum history** and note the number of the last **yum** command you used.
7. Type **yum history undo *<nn>*** (where *<nn>* is replaced with the number that you found in step 6). This undoes the last action, so it removes the package you just installed.
8. Repeat the **yum list setools-console** command. The package is now listed as available but not as installed.

Managing Package Module Streams

In RHEL 8, Red Hat has introduced Package Module Streams. In previous versions of RHEL, it was difficult to work with different versions of packages. For instance, RHEL 7 has always provided Python 2.x as the default Python version, even though Python 3.x was already available. Using a different version of user-space packages like Python was very difficult, as different repositories had to be added to offer Python 3 instead, which on its turn would again create problems for users who would rather use an older version of Python. In RHEL 8 different versions of the same package can be offered using Package Module Streams. Also, Package Module Streams make it possible to update critical applications more frequently than the base operating system is updated.

To separate core operating system packages from user-space packages, RHEL 8 provides two main repositories, referenced earlier: BaseOS and Application Stream. In BaseOS you'll find core operating system packages as RPM packages. The life cycle of these packages is comparable to the life cycle of previous RHEL versions.

Understanding Modules

In the Application Stream repository, content with varying life cycles is provided. This content is provided as the traditional RPM packages, but also as modules. A *module* describes a set of RPM packages that belong together. Typically, modules are organized around a specific version of an application, and in a module you'll find module packages, together with all of the dependencies for that specific version.

Each module can have one or more application streams. A *stream* contains one specific version, and updates are provided for a specific stream. When working with modules that have different streams, only one stream can be enabled at the same time.

Modules can also have one or more profiles. A *profile* is a list of packages that are installed together for a particular use case. You may find, for instance, a minimal profile, a default profile, a server profile. and many more. While working with modules, you may select which profile you want to use. Table 9-4 provides an overview of key terminology when working with modules.

Table 9-4 YUM Module Terminology

Item	Explanation
RPM	The default package format. Contains files, as well as metadata that describes how to install the files. Optionally may contain pre- and post-installation scripts as well.
Module	A delivery mechanism to install RPM packages. In a module different versions and profiles can be provided.
Application stream	A specific version of the module.
Profile	A collection of packages that are installed together for a particular use case.

Managing Modules

The **yum** command in RHEL 8 supports working with modules using the **yum module** command. To find out which modules are available, you may want to start with the **yum module list** command. You can see its sample output in Example 9-13.

Example 9-13 Showing Yum Modules with **yum module list**

```
[root@server3 ~]# yum module list
Updating Subscription Management repositories.
Unable to read consumer identity
This system is not registered to Red Hat Subscription Management.
You can use subscription-manager to register.
Last metadata expiration check: 0:21:12 ago on Sat 25 May 2019
09:36:47 AM EDT.

AppStream
Name                               Stream      Profiles    Summary
389-ds                             1.4         common [d]  389 Directory Server
                                     (base)
ant                                1.10 [d]    common [d]  Java build tool
container-tools                    1.0         common [d]  Common tools and
                                     dependencies for
                                     container runtimes
container-tools                    rhel8 [d][e] common [d]  Common tools and
                                     dependencies for
                                     container runtimes
freeradius                         3.0 [d]     server [d]  High-performance and
                                     highly configurable
                                     free RADIUS server
gimp                               2.8 [d]     common [d], devel gimp module
go-toolset                         rhel8 [d]    common [d]  Go
httpd                              2.4 [d]     common [d], devel Apache HTTP Server
                                     , minimal
idm                                DL1         common [d], adtru The Red Hat
                                     Enterprise Linux Identity
                                     Managest, client, dns, ment
                                     system module server
idm                                client [d]   common [d]  RHEL IdM long term
                                     support client module
inkscape                           0.92.3 [d]  common [d]  Vector-based drawing
                                     program using SVG
javapackages-runtime              201801 [d]  common [d]  Basic runtime
                                     utilities to
                                     support Java
                                     applications
...
pki-core                           10.6        PKI Core
pki-deps                           10.6        PKI Dependencies
```


postgresql	10 [d]	client, server [d]	PostgreSQL server and client module
postgresql	9.6	client, server [d]	PostgreSQL server and client module
python27	2.7 [d]	common [d]	Python programming language, version 2.7
python36	3.6 [d][e]	common [d], build	Python programming language, version 3.6
redis	5 [d]	common [d]	Redis persistent key-value database
rhn-tools	1.0 [d]	common [d]	Red Hat Satellite 5 tools for RHEL
ruby	2.5 [d]	common [d]	An interpreter of object-oriented scripting language
rust-toolset	rhel8 [d]	common [d]	Rust
satellite-5-client	1.0 [d][e]	common [d], gui	Red Hat Satellite 5 client packages
scala	2.10 [d]	common [d]	A hybrid functional/object-oriented language for the JVM
squid	4 [d]	common [d]	Squid - Optimising Web Delivery
subversion	1.10 [d]	common [d], server	Apache Subversion
swig	3.0 [d]	common [d], complete	Connects C/C++/Objective C to some high-level programming languages
varnish	6 [d]	common [d]	Varnish HTTP cache
virt	rhel [d][e]	common [d]	Virtualization module

Hint: [d]efault, [e]nabled, [x]disabled, [i]nstalled

In the list of modules, you can see whether or not the module is installed and whether or not a specific stream is enabled. To list specific streams for a module, use the **yum module list *modulename*** command. For instance, use **yum module list perl** to get details about streams that are available for the Perl module, as shown in Example 9-14.

Example 9-14 Showing Details About YUM Modules with **yum module list**

```
[root@server3 ~]# yum module list perl
Updating Subscription Management repositories.
Unable to read consumer identity
This system is not registered to Red Hat Subscription Management.
You can use subscription-manager to register.
Last metadata expiration check: 0:26:17 ago on Sat 25 May 2019
09:36:47 AM EDT.

AppStream
Name           Stream           Profiles          Summary
perl           5.24             common [d], minimal Practical
Extraction and Report Language
perl           5.26 [d]         common [d], minimal Practical
Extraction and Report Language
Hint: [d]efault, [e]nabled, [x]disabled, [i]nstalled
```

After finding out which module streams are available, the next step is to get information about specific profiles. You can use **yum module info** to obtain this information. For instance, use **yum module info perl** to get more information about the Perl module. This will provide information for profiles that are available in all the module streams. To find profile information for a specific stream, you can provide the stream version as an argument. For instance, use **yum module info perl:5.26** (see Example 9-15).

Example 9-15 Showing Information About YUM Modules with **yum module list**

```
[root@server3 ~]# yum module info perl:5.26
Updating Subscription Management repositories.
Unable to read consumer identity
This system is not registered to Red Hat Subscription Management.
You can use subscription-manager to register.
Last metadata expiration check: 0:29:59 ago on Sat 25 May 2019
09:36:47 AM EDT.

Name           : perl
Stream          : 5.26 [d][a]
Version         : 820181219174508
Context        : 9edba152
Profiles       : common [d], minimal
Default profiles : common
Repo           : AppStream
Summary        : Practical Extraction and Report Language
```

Description : Perl is a high-level programming language with roots in C, sed, awk and shell scripting. Perl is good at handling processes and files, and is especially good at handling text. Perl's hallmarks are practicality and efficiency. While it is used to do a lot of different things, Perl's most common applications are system administration utilities and web programming.

Hint: [d]efault, [e]nabled, [x]disabled, [i]nstalled, [a]ctive]

After finding module information, the next step is to enable a module stream and install modules. Every module has a default module stream, providing access to a specific version. If that version is what you need, you don't have to enable anything. If you want to work with a different version, you should start by enabling the corresponding module stream. For example, type **yum module enable perl:5.24** to enable that specific version.

Enabling a module stream before starting to work with a specific module is not mandatory. If you just use **yum module install** to install packages from a module, packages from the default module stream will be installed. You can also switch between application stream versions. If, for instance, you are now on php:7.1 and you want to change to php:7.2, you just have to type **yum module install php:7.2**. This will disable the old stream and enable the new stream. After doing this, to ensure that all dependent packages that are not in the module itself are updated as well, you should type **yum distro-sync** to finalize the procedure.

In Exercise 9-3 you can practice working with Yum modules in RHEL 8.

Exercise 9-3 Working with Modules

1. Type **yum module list**. You'll see a list of all modules that are available.
2. Type **yum module provides httpd**. This will show you that the httpd package is a part of the httpd module. As this module has just one stream, we'll have a look at the PHP module in the next steps in this exercise.
3. Type **yum module info php** to get information about the php module. This gives information about the default PHP application stream.
4. To investigate packages in the application stream for PHP 7.1, type **yum module info --profile 7.1**.

5. Type **yum module install php:7.1/devel**. This will install PHP from the PHP 7.1 application stream, using the developer profile. Notice that by installing it from this specific stream, you'll also enable the 7.1 application stream.
6. Type **yum module install php:7.2** to switch to the newer version of the PHP module.
7. To ensure that all dependent packages are updated as well, type **yum distro-sync**.

Managing Software Packages with rpm

Once upon a time, repositories did not exist, and the **rpm** command was used to install package files after they had been downloaded. That worked, but there was one major issue: the *dependency hell*. Because RPM packages have always focused on specific functionality, to install specific software, a collection of RPM packages was normally required. Therefore, a “missing dependency” message was often issued while users were trying to install RPM packages, which meant that to install the selected package, other packages needed to be installed first. Sometimes a whole chain of dependencies needed to be installed to finally get the desired functionality. That did not make working with RPM packages a joyful experience.

On modern RHEL systems, repositories are used, and packages are installed using **yum**. The **yum** command considers all package dependencies and tries to look them up in the currently available repositories. On a RHEL system configured to get updates from RHN, or on a CentOS system where consistent repositories are used, the result is that package installation nowadays is without problems and the **rpm** command no longer is needed for software installation.

Even after downloading an RPM package file, you do not need to use the **rpm -Uvh packagename** command anymore to install it (even if it still works). A much better alternative is **yum install packagename**, which installs the package and also considers the repositories to resolve dependencies automatically. (In earlier versions of RHEL, the **yum localinstall** command was used to do this; in RHEL 8, **yum localinstall** is deprecated.) That does not mean the **rpm** command has become totally useless. You can still use it to query RPM packages.

TIP On your system, two package databases are maintained: the YUM database and the RPM database. When you are installing packages through **yum**, the YUM database is updated first, after which the updated information is synchronized to the RPM database. If you install packages using the **rpm** command, the update is written to the RPM database only and will not be updated to the YUM database, which is an important reason not to use the **rpm** command anymore to install software packages.

Understanding RPM Filenames

When working with RPM packages directly, it makes sense to understand how the RPM filename is composed. A typical RPM filename looks like `autofs-5.0.7-40.el7.x86_64.rpm`. This name consists of several parts:

Key Topic

- **autofs:** The name of the actual package.
- **5.0.7:** The version of the package. This normally corresponds to the name of the package as it was released by the package creator.
- **-40:** The subversion of the package.
- **el7:** The Red Hat version this package was created for.
- **x86_64:** The platform (32 bits or 64 bits) this package was created for.

Querying the RPM Database

The **rpm** command enables you to get much information about packages. Using RPM queries can be a really useful way to find out how software can be configured and used. To start, you can use the **rpm -qa** command. Like **yum list installed**, this shows a list of all software that is installed on the machine. Use **grep** on this command to find out specific package names. To perform queries on RPM packages, you just need the name and not the version information.

After finding the package about which you want to have more information, you can start with some generic queries to find out what is in the package. In the following examples, I assume that you are using RPM queries on the `nmap` RPM package. To start, type **rpm -qi nmap** to get a description of the package. This will perform a query of a package that is already installed on your system, and it will query the package database to get more details about it.

The next step is to use **rpm -ql nmap**, which shows a list of all files that are in the package. On some packages, the result can be a really long list of filenames that is not particularly useful. To get more specific information, use **rpm -qd nmap**, which shows all documentation available for the package, or **rpm -qc nmap**, which shows all configuration files in the package.

Using RPM queries can really help in finding out more useful information about packages. The only thing that you need to know is the RPM package name that a specific file belongs to. To find this, use **rpm -qf**, followed by the specific filename you are looking for. Use, for instance, **rpm -qf /bin/ls** to find the name of the RPM package the **ls** command comes from. In upcoming Exercise 9-4, you'll see how useful it can be to use RPM queries in this way.

Querying RPM Package Files

RPM queries by default are used on the RPM database, and what you are querying are installed RPM packages. It sometimes makes sense to query an RPM package file before actually installing it. To do this, you need to add the **-p** option to the query, because without the **-p** option, you will be querying the database, not the package file. Also, when querying a package file, you need to refer to the complete filename, including the version number and all other information that you do not have to use when querying the RPM database. As an example, the **rpm -qp --scripts httpd-2.4.6-19.el7.centos.x86_64.rpm** command queries the specific RPM file to see whether it contains scripts.

A query option that needs special attention is **--scripts**, which queries an RPM package or package file to see which scripts it contains (if any). This option is especially important when combined with the **-p** option, to find out whether a package file that you are going to install includes any scripts.

When you install RPM packages, you do so as root. Before installing an RPM package from an unknown source, you need to make sure that it does not include any rogue scripts. If you do not, you risk installing malware on your computer without even knowing it.

Table 9-5 describes the most important RPM querying options.

Table 9-5 Common RPM Query Commands

Command	Use
rpm -qf	Uses a filename as its argument to find the specific RPM package a file belongs to.
rpm -ql	Uses the RPM database to provide a list of files in the RPM package.
rpm -qi	Uses the RPM database to provide package information (equivalent to yum info).
rpm -qd	Uses the RPM database to show all documentation that is available in the package.
rpm -qc	Uses the RPM database to show all configuration files that are available in the package.
rpm -q --scripts	Uses the RPM database to show scripts that are used in the package. This is particularly useful if combined with the -p option.
rpm -qp <pkg>	The -p option is used with all the previously listed options to query individual RPM package files instead of the RPM package database. Using this option before installation helps you find out what is actually in the package before it is installed.
rpm -qR	Shows dependencies for a specific package.

**Key
Topic**

Command	Use
rpm -V	Shows which parts of a specific package have been changed since installation.
rpm -Va	Verifies all installed packages and shows which parts of the package have been changed since installation. This is an easy and convenient way to do a package integrity check.
rpm -qa	Lists all packages that are installed on this server.

Using repoquery

While **rpm -qp** provides useful tools to query packages before installation, there is a slight problem with this command: It works only on RPM package files, and it cannot query files directly from the repositories. If you want to query packages from the repositories before they have been installed, you need **repoquery**. This binary is not installed by default, so make sure to install the **yum-utils** RPM package to use it.

The **repoquery** command is pretty similar to the **rpm -q** command and uses many similar options. There is just one significant option missing: **--script**. A simple solution is to make sure that you are using trusted repositories only, to prevent installing software that contains dangerous script code.

If you need to thoroughly analyze what an RPM package is doing when it is installed, you can download it to your machine, which allows you to use the **rpm -qp --scripts** command on the package. To download a package from the repository to the local directory, you can use the **yumdownloader** command, which comes from the **yum-utils** package.

Now that you have learned all about RPM querying options, you can practice these newly acquired skills in Exercise 9-4 to get more information about software that is installed on your RHEL system.

Exercise 9-4 Using RPM Queries

1. Type **which dnsmasq**. This command gives the complete path name of the **dnsmasq** command.
2. Type **rpm -qf \$(which dnsmasq)**. This will do an RPM file query on the result of the **which dnsmasq** command; you learn more about this technique in Chapter 19, “An Introduction to Bash Shell Scripting.”
3. Now that you know that the **dnsmasq** binary comes from the **dnsmasq** package, use **rpm -qi dnsmasq** to show more information about the package.

4. The information that is shown with **rpm -qi** is useful, but it does not give the details that are needed to start working with the software in the package. Use **rpm -ql dnsmasq** to show a list of all files in the package.
5. Use **rpm -qd dnsmasq** to show the available documentation. Notice that this command reveals that there is a man page, but there is also a doc.html file and a setup.html file in the /usr/share/doc/dnsmasq-version directory. Open these files with your browser to get more information about the use of dnsmasq.
6. Type **rpm -qc dnsmasq** to see which configuration files are used by dnsmasq.
7. After installation, it does not make much sense, but it is always good to know which scripts are executed when a package is installed. Use **rpm -q --scripts dnsmasq** to show the script code that can be executed from this RPM.

TIP Working with RPM queries is an extremely valuable skill on the RHCSA exam. If you know how to handle queries, you can find all relevant configuration files and the documentation.

Summary

In this chapter, you learned how to work with software on Red Hat Enterprise Linux. You learned how to use **yum** to manage software packages coming from repositories. You also learned how to use the **rpm** command to perform queries on the packages on your system. Make sure that you master these essential skills well; they are key to getting things done on Red Hat Enterprise Linux.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 9-6 lists a reference of these key topics and the page number on which each is found.

**Key
Topic****Table 9-6** Key Topics for Chapter 9

Key Topic Element	Description	Page
Paragraph	Description of how to create a repository	203
Table 9-3	Common yum tasks	210
List	RPM package name components	226
Table 9-5	Common RPM query commands	227

Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

YUM, repository, dependency, package, Red Hat Network (RHN), package groups, dependency hell, RPM, module, application stream, application profile

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. You have a directory containing a collection of RPM packages and want to make that directory a repository. Which command enables you to do that?
2. What needs to be in the repository file to point to a repository on `http://server.example.com/repo`?
3. You have just configured a new repository to be used on your RHEL computer. Which command enables you to verify that the repository is indeed available?
4. Which command enables you to search the RPM package containing the file `useradd`?
5. Which two commands do you need to use to show the name of the yum group that contains security tools and shows what is in that group?
6. Which command do you use to ensure that all PHP-related packages are going to be installed using the older version 5.1, without actually installing anything yet?

7. You want to make sure that an RPM package that you have downloaded does not contain any dangerous script code. Which command enables you to do so?
8. Which command reveals all documentation in an RPM package?
9. Which command shows the RPM package a file comes from?
10. Which command enables you to query software from the repository?

End-of-Chapter Labs

In these end-of-chapter labs, you use some of the essential RHEL package management skills. All assignments can be done on one server.

Lab 9.1

1. Copy some RPM files from the installation disk to the /myrepo directory. Make this directory a repository and make sure that your server is using this repository.
2. List the repositories currently in use on your server.
3. Search for the package that contains the cache-only DNS name server. Do not install it yet.
4. Perform an extensive query of the package so that you know before you install it which files it contains, which dependencies it has, and where to find the documentation and configuration.
5. Check whether the RPM package contains any scripts. You may download it, but you may not install it yet; you want to know which scripts are in a package before actually installing it, right?
6. Install the package you have found in step 3.
7. Undo the installation.

Lab 9.2

1. Find out which application streams are available for the Perl module.
2. Install Perl from the latest application stream version, using the most minimal installation option.
3. Downgrade, such that the previous version of the Perl application will be used from now on.
4. Ensure that all dependent packages on your system are downgraded as well.



The following topics are covered in this chapter:

- Introduction to Process Management
- Managing Shell Jobs
- Using Common Command-Line Tools for Process Management
- Using **top** to Manage Processes
- Using **tuned** to Optimize Performance

The following RHCSA exam objectives are covered in this chapter:

- Identify CPU/memory-intensive processes and kill processes
- Adjust process scheduling

Process management is an important task for a Linux administrator. In this chapter, you learn what you need to know to manage processes from a perspective of daily operation of a server. You'll learn how to work with shell jobs and generic processes. You'll also be introduced to system performance optimization using **tuned**.

Managing Processes

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 10-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 10-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Introduction to Process Management	1
Managing Shell Jobs	2–3
Using Common Command-Line Tools for Process Management	4–8
Using top to Manage Processes	9
Using tuned to Optimize Performance	10

1. Which of the following is not generally considered a type of process? (Choose two.)
 - a. A shell job
 - b. A cron job
 - c. A daemon
 - d. A thread
2. Which of the following can be used to move a job to the background?
 - a. Press **&**
 - b. Press Ctrl-Z and then type **bg**
 - c. Press Ctrl-D and then type **bg**
 - d. Press Ctrl-Z, followed by **&**

3. Which key combination enables you to cancel a current interactive shell job?
 - a. Ctrl-C
 - b. Ctrl-D
 - c. Ctrl-Z
 - d. Ctrl-Break
4. Which of the following statements are true about threads? (Choose two.)
 - a. Threads cannot be managed individually by an administrator.
 - b. Multithreaded processes can make the working of processes more efficient.
 - c. Threads can be used only on supported platforms.
 - d. Using multiple processes is more efficient, in general, than using multiple threads.
5. Which of the following commands is most appropriate if you're looking for detailed information about the command and how it was started?
 - a. **ps ef**
 - b. **ps aux**
 - c. **ps**
 - d. **ps fax**
6. Of the following **nice** values, which will increase the priority of the selected process?
 - a. 100
 - b. 20
 - c. -19
 - d. -100
7. Which of the following shows correct syntax to change the priority for the current process with PID 1234?
 - a. **nice -n 5 1234**
 - b. **renice 5 1234**
 - c. **renice 5 -p 1234**
 - d. **nice 5 -p 1234**

8. Which of the following commands cannot be used to send signals to processes?
 - a. **kill**
 - b. **killall**
 - c. **pkill**
 - d. **killall**

9. Which of the following commands would you use from **top** to change the priority of a process?
 - a. **r**
 - b. **n**
 - c. **c**
 - d. **k**

10. Which of the following commands will set the current performance profile to powersave?
 - a. **tuneadm profile set powersave**
 - b. **tuned-adm profile powersave**
 - c. **tuneadm profile --set powersave**
 - d. **tuned-adm profile --set powersave**

Foundation Topics

Introduction to Process Management

For everything that happens on a Linux server, a process is started. For that reason, process management is among the key skills that an administrator has to master. To do this efficiently, it is important to know which type of process you are dealing with. A major distinction can be made between three process types:

- *Shell jobs* are commands started from the command line. They are associated with the shell that was current when the process was started. Shell jobs are also referred to as *interactive processes*.
- *Daemons* are processes that provide services. They normally are started when a computer is booted and often (but certainly not in all cases) run with root privileges.
- *Kernel threads* are a part of the Linux kernel. You cannot manage them using common tools, but for monitoring of performance on a system, it's important to keep an eye on them.

When a process is started, it can use multiple threads. A *thread* is a task started by a process and that a dedicated CPU can service. The Linux shell does not offer tools to manage individual threads. Thread management should be taken care of from within the command.

To manage a process efficiently, it is paramount that you know what type of process you are dealing with. Shell jobs require a different approach than the processes that are automatically started when a computer boots.

Managing Shell Jobs

When a user types a command, a shell job is started. If no particular measures have been taken, the job is started as a foreground process, occupying the terminal it was started from until it has finished its work. As a Linux administrator, you need to know how to start shell jobs in the foreground or background and what you can do to manage shell jobs.

Running Jobs in the Foreground and Background

By default, any executed command is started as a foreground job. That means that you cannot do anything on the terminal where the command was started until it is done. For many commands, that does not really matter because the command often takes a little while to complete, after which it returns access to the shell from

which it was started. Sometimes it might prove useful to start commands in the background. This makes sense for processes that do not require user interaction and take significant time to finish. A process that does require user interaction will not be able to get that when running in the background, and for that reason will typically stall when moved to the background. You can take two different approaches to run a process in the background.

If you know that a job will take a long time to complete, you can start it with an **&** behind it. This immediately starts the job in the background to make room for other tasks to be started from the command line. To move the last job that was started in the background back as a foreground job, use the **fg** command. This command immediately, and with no further questions, brings the last job back to the foreground. If multiple jobs are currently running in the background, you can move a job back to the foreground by adding its job ID, as shown by the **jobs** command.

A job might sometimes have been started that takes (much) longer than predicted. If that happens, you can use Ctrl-Z to temporarily stop the job. This does not remove the job from memory; it just pauses the job so that it can be managed. Once paused, you can continue it as a background job by using the **bg** command. An alternative key sequence that you can use to manage shell jobs is Ctrl-C. This stops the current job and removes it from memory.

A related key combination is Ctrl-D, which sends the End Of File (EOF) character to the current job. The result is that the job stops waiting for further input so that it can complete what it was currently doing. The result of pressing Ctrl-D is sometimes similar to the result of pressing Ctrl-C, but there is a difference. When Ctrl-C is used, the job is just canceled, and nothing is closed properly. When Ctrl-D is used, the job stops waiting for further input and next terminates, which often is just what is needed to complete in a proper way.

Managing Shell Jobs

When moving jobs between the foreground and background, it may be useful to have an overview of all current jobs. To get such an overview, use the **jobs** command. As you can see in Table 10-2, this command gives an overview of all jobs currently running as a background job, including the job number assigned to the job when starting it in the background. These job numbers can be used as an argument to the **fg** and **bg** commands to perform job management tasks. In Exercise 10-1, you learn how to perform common job management tasks from the shell.

**Key
Topic****Table 10-2** Job Management Overview

Command	Use
& (used at the end of a command line)	Starts the command immediately in the background.
Ctrl-Z	Stops the job temporarily so that it can be managed. For instance, it can be moved to the background.
Ctrl-D	Sends the End Of File (EOF) character to the current job to indicate that it should stop waiting for further input.
Ctrl-C	Can be used to cancel the current interactive job.
bg	Continues the job that has just been frozen using Ctrl-Z in the background.
fg	Brings back to the foreground the last job that was moved to background execution.
jobs	Shows which jobs are currently running from this shell. Displays job numbers that can be used as an argument to the commands bg and fg .

Exercise 10-1 Managing Jobs

1. Open a root shell and type the following commands:

```
sleep 3600 &
dd if=/dev/zero of=/dev/null &
sleep 7200
```
2. Because you started the last command with no **&** after the command, you have to wait 2 hours before you get control of the shell back. Press **Ctrl-Z** to stop the command.
3. Type **jobs**. You will see the three jobs that you just started. The first two of them have the Running state, and the last job currently is in the Stopped state.
4. Type **bg 3** to continue running job 3 in the background. Note that because it was started as the last job, you did not really have to add the number 3.
5. Type **fg 1** to move job 1 to the foreground.
6. Press **Ctrl-C** to cancel job number 1 and type **jobs** to confirm that it is now gone.
7. Use the same approach to cancel jobs 2 and 3 also.
8. Open a second terminal on your server.
9. From that second terminal, type **dd if=/dev/zero of=/dev/null &**.
10. Type **exit** to close the second terminal.
11. From the other terminal, start **top**. You will see that the **dd** job is still running. From **top**, use **k** to kill the **dd** job.

NOTE You learned how to manage interactive shell jobs in this section. Note that all of these jobs are processes as well. As the user who started the job, you can also manage it. In the next section, you learn how to use process management to manage jobs started by other users.

Managing Parent-Child Relations

When a process is started from a shell, it becomes a child process of that shell. In process management, the parent-child relationship between processes is very important. The parent is needed to manage the child. For that reason, all processes started from a shell are terminated when that shell is stopped. This also offers an easy way to terminate processes no longer needed.

Processes started in the background will not be killed when the parent shell from which they were started is killed. To terminate these processes, you need to use the **kill** command, as described later in this chapter.

NOTE In earlier versions of the Bash shell, background processes were also killed when the shell they were started from was terminated. To prevent that, the process could be started with the **nohup** command in front of it. Using **nohup** for this purpose is no longer needed in RHEL 8. If a parent process is killed while the child process still is active, the child process becomes a child of **systemd** instead.

Using Common Command-Line Tools for Process Management

On a Linux server, many processes are usually running. On an average server or desktop computer, there are often more than 100 active processes. With so many processes being active, things may go wrong. If that happens, it is good to know how noninteractive processes can be stopped or how the priority of these processes can be adjusted to make more system resources available for other processes.

Understanding Processes and Threads

Tasks on Linux are typically started as processes. One process can start several worker threads. Working with threads makes sense, because if the process is very busy, the threads can be handled by different CPUs or CPU cores available in the machine. As a Linux administrator, you cannot manage individual threads; you can manage processes, though. It is the programmer of the multithreaded application that has to define how threads relate to one another.

Before talking about different ways to manage processes, it is good to know that there are two different types of background processes: kernel threads and daemon processes. *Kernel threads* are a part of the Linux kernel, and each of them is started with its own process identification number (PID). When managing processes, it is easy to recognize the kernel processes because they have a name that is between square brackets. Example 10-1 shows a list of a few processes as output of the command **ps aux | head** (discussed later in this chapter), in which you can see a couple of kernel threads.

As an administrator, it is important to know that kernel threads cannot be managed. You cannot adjust their priority; neither is it possible to kill them, except by taking the entire machine down.

Example 10-1 Showing Kernel Threads with **ps aux**

[root@server3 ~]# ps aux head										
USER	PID	%CPU	%MEM	VSZ	RSS	TTY	STAT	START	TIME	COMMAND
root	1	0.0	0.4	252864	7792	?	Ss	08:25	0:02	/usr/lib/systemd/systemd --switched-root --system --deserialize 17
root	2	0.0	0.0	0	0	?	S	08:25	0:00	[kthreadd]
root	3	0.0	0.0	0	0	?	I<	08:25	0:00	[rcu_gp]
root	4	0.0	0.0	0	0	?	I<	08:25	0:00	[rcu_par_gp]
root	6	0.0	0.0	0	0	?	I<	08:25	0:00	[kworker/0:0H-kblockd]
root	8	0.0	0.0	0	0	?	I<	08:25	0:00	[mm_percpu_wq]
root	9	0.0	0.0	0	0	?	S	08:25	0:00	[ksoftirqd/0]
root	10	0.0	0.0	0	0	?	I	08:25	0:00	[rcu_sched]
root	11	0.0	0.0	0	0	?	S	08:25	0:00	[migration/0]

Using ps to Get Process Information

The most common command to get an overview of currently running processes is **ps**. If used without any arguments, the **ps** command shows only those processes that have been started by the current user. You can use many different options to display different process properties. If you are looking for a short summary of the active processes, use **ps aux** (as you saw in Example 10-1). If you are looking for not only the name of the process but also the exact command that was used to start the process, use **ps -ef** (see Example 10-2). Alternative ways to use **ps** exist as well, such as the command **ps fax**, which shows hierarchical relationships between parent and child processes (see Example 10-3).

Example 10-2 Using **ps -ef** to See the Exact Command Used to Start Processes

```
[root@server3 ~]# ps -ef
```

UID	PID	PPID	C	STIME	TTY	TIME	CMD
root	1	0	0	08:25	?	00:00:02	/usr/lib/systemd/systemd --switched-root --system --deserialize 17
...							
root	34948	2	0	12:16	?	00:00:00	[kworker/0:1-events]
root	34971	1030	0	12:17	?	00:00:00	sshd: root [priv]
root	34975	34971	0	12:17	?	00:00:00	sshd: root@pts/2
root	34976	34975	0	12:17	pts/2	00:00:00	-bash
root	35034	1	0	12:17	pts/2	00:00:00	sleep 3600
root	35062	2	0	12:20	?	00:00:00	[kworker/u256:2]
root	35064	2	0	12:20	?	00:00:00	[kworker/0:3-cgroup_ destroy]
root	35067	2	0	12:20	?	00:00:00	[kworker/1:2-events_ freezable_power_]
root	35087	939	0	12:21	?	00:00:00	sleep 60
root	35088	33127	0	12:22	pts/1	00:00:00	ps -ef

NOTE For many commands, options need to start with a hyphen. For some commands, this is not the case and using the hyphen is optional. The **ps** command is one of these commands, due to historic reasons. In the old times of UNIX, there were two main flavors: the System V flavor, in which using hyphens before options was mandatory, and the BSD flavor, in which using hyphens was optional. The **ps** command is based on both of these flavors, and for that reason some options don't have to start with a hyphen.

Example 10-3 Using **ps fax** to Show Parent-Child Relationships Between Processes

```
[root@server3 ~]# ps fax
```

PID	TTY	STAT	TIME	COMMAND
2	?	S	0:00	[kthreadd]
3	?	I<	0:00	_ [rcu_gp]
4	?	I<	0:00	_ [rcu_par_gp]
...				
2460	?	Ssl	0:00	_ /usr/bin/pulseaudio --daemonize=no
2465	?	Ssl	0:00	_ /usr/bin/dbus-daemon --session --address=systemd: --nofork --nopidfile --systemd-activation --

2561	?	Ssl	0:00	_ /usr/libexec/at-spi-bus-launcher
2566	?	Sl	0:00	_ /usr/bin/dbus-daemon --config-file=/usr/share/defaults/at-spi2/accessibility.conf --nofork
2569	?	Sl	0:00	_ /usr/libexec/at-spi2-registryd --use-gnome-session
2589	?	Ssl	0:00	_ /usr/libexec/xdg-permission-store
2594	?	Sl	0:00	_ /usr/libexec/ibus-portal
2704	?	Sl	0:00	_ /usr/libexec/dconf-service
2587	?	Sl	0:00	/usr/libexec/ibus-x11 --kill-daemon
2758	?	Sl	0:00	/usr/bin/gnome-keyring-daemon --daemonize --login
2908	tty3	Sl	0:00	/usr/libexec/ibus-x11 --kill-daemon
2936	?	Ssl	0:00	/usr/libexec/geoclue
3102	tty3	Sl+	0:00	/usr/libexec/gsd-printer
3173	tty3	Sl+	0:12	/usr/bin/vmtoolsd -n vmusr
3378	?	Ssl	0:00	/usr/libexec/fwupd/fwupd
3440	?	Ss	0:00	gpg-agent --homedir /var/lib/fwupd/gnupg --use-standard-socket --daemon
3455	?	S	0:00	/usr/libexec/platform-python /usr/libexec/rhsmc
33093	?	Ss	0:00	/usr/lib/systemd/systemd --user
33105	?	S	0:00	_ (sd-pam)
33117	?	S<sl	0:00	_ /usr/bin/pulseaudio --daemonize=no
33123	?	Ssl	0:00	_ /usr/bin/dbus-daemon --session --address=systemd: --nofork --nopidfile --systemd-activation --
35034	pts/2	S	0:00	sleep 3600

An important piece of information to get out of the **ps** command is the PID. Many tasks require the PID to operate, and that is why a command like **ps aux | grep dd**, which will show process details about **dd**, including its PID, is quite common. An alternative way to get the same result is to use the **pgrep** command. Use **pgrep dd** to get a list of all PIDs that have a name containing the string "dd".

Adjusting Process Priority with nice

When Linux processes are started, they are started with a specific priority. By default, all regular processes are equal and are started with the same priority, which is the priority number 20, as shown by utilities like **top**. In some cases, it is useful to change the default priority that was assigned to the process when it was started. You can do that using the **nice** and **renice** commands. Use **nice** if you want to start

a process with an adjusted priority. Use **renice** to change the priority for a currently active process. Alternatively, you can use the **r** command from the **top** utility to change the priority of a currently running process.

Changing process priority may make sense in two different scenarios. Suppose, for example, that you are about to start a backup job that does not necessarily have to finish fast. Typically, backup jobs are rather resource intensive, so you might want to start the backup job in a way that does not annoy other users too much, by lowering its priority.

Another example is where you are about to start a very important calculation job. To ensure that it is handled as fast as possible, you might want to give it an increased priority, taking away CPU time from other processes.

On earlier Linux versions, it could be dangerous to increase the priority of one job too much, because of the risk that other processes (including vital kernel processes) might be blocked out completely. On current Linux kernels, that risk is minimized for these reasons:

- Modern Linux kernels differentiate between essential kernel threads that are started as real-time processes and normal user processes. Increasing the priority of a user process will never be able to block out kernel threads or other processes that were started as real-time processes.
- Modern computers often have multiple CPU cores. A single-threaded process that is running with the highest priority will never be able to get beyond the boundaries of the CPU it is running on.

When using **nice** or **renice** to adjust process priority, you can select from values ranging from -20 to 19 . The default niceness of a process is set to 0 (which results in the priority value of 20). By applying a negative niceness, you increase the priority. Use a positive niceness to decrease the priority. It is a good idea not to use the ultimate values immediately. Instead, use increments of 5 and see how it affects the application.

TIP Do not set process priority to -20 ; it risks blocking other processes from getting served.

Let's take a look at examples of how to use **nice** and **renice**. The command **nice -n 5 dd if=/dev/zero of=/dev/null &** starts an infinite I/O-intensive job, but with an adjusted niceness so that some room remains for other processes as well. To adjust the niceness of a currently running process, you need the PID of that process. The following two commands show how **ps aux** is used to find the PID of the **dd**

job from the previous example. Next, you see how the **renice** command is used to change the niceness of that command:

1. Use **ps aux | grep dd** to find the PID of the **dd** command that you just started. The PID is in the second column of the command output.
2. Use **renice -n 10 -p 1234** (assuming that 1234 is the PID you just found).

Note that regular users can only decrease the priority of a running process. You must be root to give processes increased priority.

Sending Signals to Processes with **kill**, **killall**, and **pkill**

Before starting to think about using the **kill** command or sending other signals to processes, it is good to know that Linux processes have a hierarchical relationship. Every process has a parent process, and as long as it lives, the parent process is responsible for the child processes it has created. In older versions of Linux, killing a parent process would also kill all of its child processes. In RHEL 8, if you kill a parent process, all of its child processes become children of the **systemd** process.

The Linux kernel allows many signals to be sent to processes. Use **man 7 signals** for a complete overview of all the available signals. Three of these signals work for all processes:

Key Topic

- The signal **SIGTERM** (15) is used to ask a process to stop.
- The signal **SIGKILL** (9) is used to force a process to stop.
- The **SIGHUP** (1) signal is used to hang up a process. The effect is that the process will reread its configuration files, which makes this a useful signal to use after making modifications to a process configuration file.

To send a signal to a process, you use the **kill** command. The most common use is the need to stop a process, which you can do by using the **kill** command followed by the PID of the process. This sends the **SIGTERM** signal to the process, which normally causes the process to cease its activity and close all open files.

Sometimes the **kill** command does not work because the process you want to kill can ignore it. In that case, you can use **kill -9** to send the **SIGKILL** signal to the process. Because the **SIGKILL** signal cannot be ignored, it forces the process to stop, but you also risk losing data while using this command. In general, it is a bad idea to use **kill -9**:

- You risk losing data.
- Your system may become unstable if other processes depend on the process you have just killed.

TIP Use **kill -l** to show a list of available signals that can be used with **kill**.

There are some commands that are related to **kill**: **killall** and **pkill**. The **pkill** command is a bit easier to use because it takes the name rather than the PID of the process as an argument. You can use the **killall** command if multiple processes using the same name need to be killed simultaneously.

Using **killall** was particularly common when Linux environments were multiprocessing instead of multithreading. In a multiprocessing environment where a server starts several commands, all with the same name, it is not easy to stop these commands one by one based on their individual PID. Using **killall** enables you to terminate all these processes simultaneously.

In a multithreaded environment, the urge to use **killall** is smaller. Because there is often just one process that is generating several threads, all these threads are terminated anyway by stopping the process that started them. You still can use **killall**, though, to terminate lots of processes with the same name that have been started on your server. In Exercise 10-2, you practice using **ps**, **nice**, **kill**, and related utilities to manage processes.

Exercise 10-2 Managing Processes from the Command Line

1. Open a root shell. From this shell, type **dd if=/dev/zero of=/dev/null &**. Repeat this command three times.
2. Type **ps aux | grep dd**. This shows all lines of output that have the letters *dd* in them; you will see more than just the **dd** processes, but that should not really matter. The processes you just started are listed last.
3. Use the PID of one of the **dd** processes to adjust the niceness, using **renice -n 5 <PID>**. Notice that in **top** you cannot easily get an overview of processes and their current priority.
4. Type **ps fax | grep -B5 dd**. The **-B5** option shows the matching lines, including the five lines before that. Because **ps fax** shows hierarchical relationships between processes, you should also find the shell and its PID from which all the **dd** processes were started.
5. Find the PID of the shell from which the **dd** processes were started and type **kill -9 <PID>**, replacing **<PID>** with the PID of the shell you just found. Because the **dd** processes were started as background processes, they are not killed when their parent shell is killed. Instead, they have been moved up and are now children of the **systemd** process.

Using top to Manage Processes

A convenient tool to manage processes is **top**. For common process management tasks, **top** is great because it gives an overview of the most active processes currently running (hence the name **top**). This enables you to easily find processes that might need attention. From **top**, you can also perform common process management tasks, such as adjusting the current process priority and killing processes. Figure 10-1 shows the interface that appears when you start **top**.

top - 12:32:16 up 7:17, 2 users, load average: 0.03, 0.02, 0.00										
Tasks: 383 total, 15 running, 368 sleeping, 0 stopped, 0 zombie										
%Cpu(s): 0.7 us, 0.7 sy, 0.0 ni, 97.6 id, 0.0 wa, 0.7 hi, 0.3 si, 0.0 st										
MiB Mem : 1806.1 total, 79.4 free, 1393.2 used, 333.5 buff/cache										
MiB Swap: 2048.0 total, 1642.2 free, 405.8 used. 211.8 avail Mem										
PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+ COMMAND
1006	root	20	0	226576	4408	3928	S	0.3	0.2	0:21.84 vmtoolsd
1199	root	20	0	0	0	0	R	0.3	0.0	0:09.51 kvdo0:cpuQ1
7996	student	20	0	2883328	103388	37904	S	0.3	5.6	0:18.13 gnome-shell
8542	student	20	0	535444	33096	19336	S	0.3	1.8	0:01.91 gnome-terminal-
14898	root	20	0	64124	4736	3852	R	0.3	0.3	0:00.06 top
1	root	20	0	244744	10384	7188	S	0.0	0.6	0:04.59 systemd
2	root	20	0	0	0	0	S	0.0	0.0	0:00.04 kthreadd
3	root	0	-20	0	0	0	I	0.0	0.0	0:00.00 rcu_gp
4	root	0	-20	0	0	0	I	0.0	0.0	0:00.00 rcu_par_gp
6	root	0	-20	0	0	0	I	0.0	0.0	0:00.00 kworker/0:0H-kblockd
8	root	0	-20	0	0	0	I	0.0	0.0	0:00.00 mm_percpu_wq
9	root	20	0	0	0	0	S	0.0	0.0	0:00.41 ksoftirqd/0
10	root	20	0	0	0	0	R	0.0	0.0	0:00.66 rcu_sched
11	root	rt	0	0	0	0	S	0.0	0.0	0:00.00 migration/0
12	root	rt	0	0	0	0	S	0.0	0.0	0:00.00 watchdog/0
13	root	20	0	0	0	0	S	0.0	0.0	0:00.00 cpuhp/0
15	root	20	0	0	0	0	S	0.0	0.0	0:00.00 kdevtmpfs
16	root	0	-20	0	0	0	I	0.0	0.0	0:00.00 netns
17	root	20	0	0	0	0	S	0.0	0.0	0:00.00 kauditd
18	root	20	0	0	0	0	S	0.0	0.0	0:00.04 khungtaskd
19	root	20	0	0	0	0	S	0.0	0.0	0:00.00 oom_reaper
20	root	0	-20	0	0	0	I	0.0	0.0	0:00.00 writeback
21	root	20	0	0	0	0	S	0.0	0.0	0:00.08 kcompactd0
22	root	25	5	0	0	0	S	0.0	0.0	0:00.00 ksmd
23	root	39	19	0	0	0	S	0.0	0.0	0:01.69 khugepaged
24	root	0	-20	0	0	0	I	0.0	0.0	0:00.00 crypto
25	root	0	-20	0	0	0	I	0.0	0.0	0:00.00 kintegrityd
26	root	0	-20	0	0	0	I	0.0	0.0	0:00.00 kblockd
27	root	0	-20	0	0	0	I	0.0	0.0	0:00.00 md
28	root	0	-20	0	0	0	I	0.0	0.0	0:00.00 edac-poller

FIGURE 10-1 Using **top** Makes Process Management Easy

Among the information that you can conveniently obtain from the **top** utility is the process state. Table 10-3 provides an overview of the different process states that you may observe.

Table 10-3 Linux Process States Overview

Key
Topic

State	Meaning
Running (R)	The process is currently active and using CPU time, or in the queue of runnable processes waiting to get services.
Sleeping (S)	The process is waiting for an event to complete.
Uninterruptable sleep (D)	The process is in a sleep state that cannot be stopped. This usually happens while a process is waiting for I/O.
Stopped (T)	The process has been stopped, which typically has happened to an interactive shell process, using the Ctrl-Z key sequence.
Zombie (Z)	The process has been stopped but could not be removed by its parent, which has put it in an unmanageable state.

Now that you know how to use the **kill** and **nice** commands from the command line, using the same functionality from **top** is even easier. From **top**, type **k**; **top** then prompts for the PID of the process you want to send a signal to. By default, the most active process is selected. After you enter the PID, **top** asks which signal you want to send. By default, signal 15 for SIGTERM is used. However, if you want to insist a bit more, you can type **9** for SIGKILL. Now press Enter to terminate the process.

To renice a running process from **top**, type **r**. You are first prompted for the PID of the process you want to renice. After entering the PID, you are prompted for the **nice** value you want to use. Enter a positive value to decrease process priority or a negative value to increase process priority.

Another important parameter you can get from **top** is the load average. The load average is expressed as the number of processes that are in a runnable state (R) or in a blocking state (D). Processes are in a runnable state if they currently are running, or waiting to be serviced. Processes are in a blocking state if they are waiting for I/O. The load average is shown for the last 1, 5, and 15 minutes, and you can see the current values in the upper-right corner of the **top** screen. Alternatively, you can use the **uptime** command to show current load average statistics (see Example 10-4).

Example 10-4 Using **uptime** for Information About Load Average

```
[root@server3 ~]# uptime
12:43:03 up 4:17, 3 users, load average: 4.90, 0.98, 0.19
```

As a rule of thumb, the load average should not be higher than the number of CPU cores in your system. You can find out the number of CPU cores in your system by using the **lscpu** command. If the load average over a longer period is higher than the number of CPUs in your system, you may have a performance problem. In Exercise 10-3 you investigate the load average statistics and learn how to manage load average.

Exercise 10-3 Managing Load Average

1. Open a root shell. From this shell, type **dd if=/dev/zero of=/dev/null &**. Repeat this command three times.
2. Type **top** and observe the current load average. After a few seconds, use **q** to quit **top**.
3. From the command line, type **uptime**. You should see the numbers that are shown as the load average slowly increasing.
4. Type **lscpu** and look for the number of CPU(s). Also look for the Core(s) per CPU parameter so that you can calculate the total number of CPU cores.
5. Use **killall dd** to kill all **dd** processes.

Using tuned to Optimize Performance

To offer the best possible performance right from the start, RHEL 8 comes with **tuned**. It offers a daemon that monitors system activity and provides some profiles. In the profiles, an administrator can automatically tune a system for best possible latency, throughput, or power consumption.

Based on the properties of an installed system, a **tuned** profile is selected automatically at installation, and after installation it's possible to manually change the current profile. Administrators can also change settings in a **tuned** profile. Table 10-4 gives an overview of the default profiles.

Table 10-4 Tuned Profile Overview

Profile	Use
balanced	The best compromise between power usage and performance
desktop	Based on the balanced profile, but tuned for better response to interactive applications
latency-performance	Tuned for maximum throughput
network-latency	Based on latency-performance, but with additional options to reduce network latency
network-throughput	Based on throughput-performance, optimizes older CPUs for streaming content
powersave	Tunes for maximum power saving
throughput-performance	Tunes for maximum throughput
virtual-guest	Optimizes Linux for running as a virtual machine
virtual-host	Optimizes Linux for use as a KVM host

It is relatively easy to create custom profiles. Also, when installing specific packages, profiles may be added. So you may find that some additional performance profiles exist on your server.

To manage the performance profile, the **tuned-adm** command is provided. It talks to the **tuned** daemon, so before you can use it, run **systemctl enable --now tuned** to start the **tuned** profile. Next, use **tuned-adm** to find out which profile currently is selected. For an overview of profiles available on your server, type **tuned-adm list**. To select another profile, type **tuned-adm profile profile-name**. The **tuned** service can also recommend a **tuned** profile for your system: use **tuned-adm recommend**. In Exercise 10-4 you can practice working with **tuned**.

Exercise 10-4 Using tuned

1. Use **yum -y install tuned** to ensure that **tuned** is installed. (It probably already is.)
2. Type **systemctl status tuned** to check whether **tuned** currently is running. If it is not, use **systemctl enable --now tuned**.
3. Type **tuned-adm active** to see which profile currently is used.
4. Type **tuned-adm recommend** to see which **tuned** profile is recommended.
5. To select and activate the **throughput-performance** profile, type **tuned-adm profile throughput-performance**.

Summary

Managing processes is a common task for a Linux system administrator. In this chapter, you learned how to look up specific processes and how to change their priority using **nice** and **kill**. You have also learned how to use **tuned** to select the performance profile that best matches your server's workload.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 10-5 lists a reference of these key topics and the page number on which each is found.



Table 10-5 Key Topics for Chapter 10

Key Topic Element	Description	Page
Table 10-2	Job management overview	238
List	Essential signal overview	244
Table 10-3	Linux process states overview	247

Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

job, process, background, foreground, **nice**, **kill**, signal, PID, thread, **tuned**, profile, zombie

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. Which command gives an overview of all current shell jobs?
2. How do you stop the current shell job to continue running it in the background?
3. Which keystroke combination can you use to cancel the current shell job?
4. A user is asking you to cancel one of the jobs he has started. You cannot access the shell that user currently is working from. What can you do to cancel his job anyway?

5. Which command would you use to show parent-child relationships between processes?
6. Which command enables you to change the priority of PID 1234 to a higher priority?
7. On your system, 20 **dd** processes are currently running. What is the easiest way to stop all of them?
8. Which command enables you to stop the command with the name **mycommand**?
9. Which command do you use from **top** to kill a process?
10. What is required to select a performance profile that best matches your system needs?

End-of-Chapter Lab

In the end-of-chapter lab, you apply some of the most important process management tasks. Use the tools that you find the most convenient to perform these labs.

Lab 10.1

1. Launch the command **dd if=/dev/zero of=/dev/null** three times as a background job.
2. Increase the priority of one of these commands using the **nice** value **-5**. Change the priority of the same process again, but this time use the value **-15**. Observe the difference.
3. Kill all the **dd** processes you just started.
4. Ensure that **tuned** is installed and active, and set the throughput-performance profile.



The following topics are covered in this chapter:

- Understanding Systemd
- Managing Units Through Systemd

The following RHCSA exam objective is covered in this chapter:

- Start and stop services and configure services to start automatically at boot

In this chapter, you'll learn about Systemd, which is the system and service manager used on RHEL 8. You'll read about all the things that Systemd can do, and once you have a good general understanding, you'll learn how to work with Systemd services. Systemd is also involved in booting your system in a desired state, which is called a target. That topic is covered in Chapter 17, "Managing and Understanding the Boot Procedure."

Working with Systemd

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 11-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 11-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Understanding Systemd	1–5
Managing Units Through Systemd	6–10

- Which command shows all service unit files on your system that are currently loaded?
 - `systemctl --type=service`
 - `systemctl --type=service --all`
 - `systemctl --list-services`
 - `systemctl --show-units | grep services`
- Which statement about Systemd wants is *not* true?
 - You can create wants by using the `systemctl enable` command.
 - The target to which a specific want applies is agnostic of the associated wants.
 - Wants are always administered in the `/usr/lib/systemd/system` directory.
 - Each service knows to which target its wants should be added.

3. What is the best solution to avoid conflicts between incompatible units?
 - a. Nothing; the unit files have defined for themselves which units they are not compatible with.
 - b. Disable the service using **systemctl disable**.
 - c. Unmask the service using **systemctl unmask**.
 - d. Mask the service using **systemctl mask**.
4. Which of the following is not a valid status for Systemd services?
 - a. Running(active)
 - b. Running(exited)
 - c. Running(waiting)
 - d. Running(dead)
5. Which of the following statements is *not* true about socket units?
 - a. A socket unit requires a service unit with the same name.
 - b. Socket units can listen on ports and activate services only when activity occurs on a port.
 - c. Socket units cannot contain the name of the associated binary that should be started.
 - d. Socket units may react upon path activity.
6. Which of the following is not a valid Systemd unit type?
 - a. service
 - b. udev
 - c. mount
 - d. socket
7. You want to find out which other Systemd units have dependencies to a specific unit. Which command would you use?
 - a. **systemd list-dependencies --reverse**
 - b. **systemctl list-dependencies --reverse**
 - c. **systemctl status my.unit --show-deps**
 - d. **systemd status my.unit --show-deps -r**

8. How do you change the default editor that Systemd is using to **vim**?
 - a. `export EDITOR=vim`
 - b. `export SYSTEMD_EDITOR=vim`
 - c. `export EDITOR=/bin/vim`
 - d. `export SYSTEMD_EDITOR=/bin/vim`
9. Which of the following keywords should you use to define a Systemd dependency if you want to ensure that the boot procedure doesn't fail if the dependency fails?
 - a. Required
 - b. Requisite
 - c. Before
 - d. Wants
10. Which of the following is *not* a valid command while working with units in `systemctl`?
 - a. `systemctl unit start`
 - b. `systemctl status -l unit`
 - c. `systemctl mask unit`
 - d. `systemctl disable unit`

Foundation Topics

Understanding Systemd

Systemd is the part of Red Hat Enterprise Linux 8 that is responsible for starting not only services but a variety of other items as well. In this chapter, you learn how Systemd is organized and what items are started from Systemd.

To describe it in a generic way, the Systemd System and Service Manager is used to start stuff. The stuff is referred to as *units*. Units can be many things. One of the most important unit types is the service. Typically, services are processes that provide specific functionality and allow connections from external clients coming in, such as the SSH service, the Apache web service, and many more. Apart from services, other unit types exist, such as socket, mount, and target. To display a list of available units, type **systemctl -t help** (see Example 11-1).

Key Topic

Example 11-1 Unit Types in Systemd

```
[root@server1 ~]# systemctl -t help
Available unit types:
service
socket
target
device
mount
automount
swap
timer
path
slice
scope
```

Understanding Systemd Unit Locations

The major benefit of working with Systemd, as compared to previous methods Red Hat used for managing services, is that it provides a uniform interface to start units. This interface is defined in the unit file. Unit files can occur in three locations:

- `/usr/lib/systemd/system` contains default unit files that have been installed from RPM packages. You should never edit these files directly.

- `/etc/systemd/system` contains custom unit files. It may also contain files that have been written by an administrator or generated by the **systemctl edit** command.
- `/run/systemd/system` contains unit files that have automatically been generated.

If a unit file exists in more than one of these locations, units in the `/run` directory have highest precedence and will overwrite any settings that were defined elsewhere. Units in `/etc/systemd/system` have second highest precedence, and units in `/usr/lib/systemd/system` come last.

Understanding Systemd Service Units

Probably the most important unit type is the service unit. It is used to start processes. You can start any type of process by using a service unit, including daemon processes and commands.

Example 11-2 gives an example of a service unit file, `vsftpd.service`, for the Very Secure FTP service.

Example 11-2 Example of a Service Unit File

```
[Unit]
Description=Vsftpd ftp daemon
After=network.target

[Service]
Type=forking
ExecStart=/usr/sbin/vsftpd /etc/vsftpd/vsftpd.conf

[Install]
WantedBy=multi-user.target
```

You can see from this unit file example that unit files are relatively easy to understand. Any Systemd *service* unit file consists of the following three sections (you'll find different sections in other types of unit files):

- **[Unit]** Describes the unit and defines dependencies. This section also contains the important **After** statement, and optionally the **Before** statement. These statements define dependencies between different units, and they relate to the perspective of this unit. The **Before** statement indicates that this unit should be started before the unit that is specified. The **After** statement indicates that this unit should be started after the unit that is specified.

**Key
Topic**

- **[Service]** Describes how to start and stop the service and request status installation. Normally, you can expect an `ExecStart` line, which indicates how to start the unit, or an `ExecStop` line, which indicates how to stop the unit. Note the **Type** option, which is used to specify how the process should start. The **forking** type is commonly used by daemon processes, but you can also use other types, such as **oneshot**, which will start any command from a Systemd unit. See **man 5 systemd.service** for more details.
- **[Install]** Indicates in which target this unit has to be started. The section “Understanding Systemd Target Units” a bit later in this chapter explains how to work with targets.

Understanding Systemd Mount Units

A mount unit specifies how a file system can be mounted on a specific directory. Example 11-3 shows an example of a mount unit file, `tmp.mount`.

Example 11-3 Example of a Mount Unit File

```
[Unit]
Description=Temporary Directory (/tmp)
Documentation=man:hier(7)
Documentation=https://www.freedesktop.org/wiki/Software/systemd/
APIFileSystems
ConditionPathIsSymbolicLink=!/tmp
DefaultDependencies=no
Conflicts=umount.target
Before=local-fs.target umount.target
After=swap.target

[Mount]
What=tmpfs
Where=/tmp
Type=tmpfs
Options=mode=1777,strictatime,nosuid,nodev
```

The `tmp.mount` unit file in Example 11-3 shows some interesting additional configuration options in its sections:

- **[Unit]** The **Conflicts** statement is used to list units that cannot be used together with this unit. Use this for mutually exclusive units.
- **[Mount]** Defines exactly where the mount has to be performed. You’ll recognize the arguments that are typically used in any **mount** command.

Understanding Systemd Socket Units

Another type of unit that is interesting to look at is the socket. A socket creates a method for applications to communicate with one another. A socket may be defined as a file, but also as a port on which Systemd will be listening for incoming connections. That way, a service doesn't have to run continuously, but instead will start only if a connection is coming in on the socket that is specified. Every socket needs a corresponding service file. Example 11-4 shows what the `cockpit.socket` file looks like; notice that this file requires a service file with the name `cockpit.service`.

Example 11-4 Example of a Socket Unit File

```
[Unit]
Description=Cockpit Web Service Socket
Documentation=man:cockpit-ws(8)
Wants=cockpit-motd.service

[Socket]
ListenStream=9090
ExecStartPost=/usr/share/cockpit/motd/update-motd '' localhost
ExecStartPost=/bin/ln -snf active.motd /run/cockpit/motd
ExecStopPost=/bin/ln -snf /usr/share/cockpit/motd/inactive.motd
    /run/cockpit/motd

[Install]
WantedBy=sockets.target
```

The important option in Example 11-4 is **ListenStream**. This option defines the TCP port that Systemd should be listening to for incoming connections. Sockets can also be created for UDP ports, in which case you would use **ListenDatagram** instead of **ListenStream**.

Understanding Systemd Target Units

**Key
Topic**

The unit files are used to build the functionality that is needed on your server. To make it possible to load them in the right order and at the right moment, a specific type of unit is used: the target unit. A simple definition of a target unit is “a group of units.” Some targets are used to define the state a server should be started in. As such, target units are comparable to the runlevels used in earlier versions of RHEL.

Other targets are just a group of services that makes it easy to manage not only individual units, but all the units that are required to get specific functionality. The

printer.target is an example of such a target; you can use it to easily start or stop all units that are required to provide printing functionality.

Targets by themselves can have dependencies on other targets. These dependencies are defined in the target unit. An example of such a dependency relation is the basic.target. This target defines all the units that should always be started. You can use the **systemctl list-dependencies** command for an overview of any existing dependencies.

Example 11-5 shows the definition of a target unit file, multi-user.target, which defines the normal operational state of a RHEL server.

Example 11-5 Example of a Target Unit File

```
[Unit]
Description=Multi-User System
Documentation=man:systemd.special(7)
Requires=basic.target
Conflicts=rescue.service rescue.target
After=basic.target rescue.service rescue.target
AllowIsolate=yes

[Install]
Alias=default.target
```

You can see that by itself the target unit does not contain much. It just defines what it requires and which services and targets it cannot coexist with. It also defines load ordering, by using the **After** statement in the [Unit] section. The target file does not contain information about the units that should be included; that is defined in the [Install] section of the different unit files.

When administrators use the **systemctl enable** command, to ensure that a unit is automatically started while booting, the [Install] section of that unit is considered to determine to which target the unit should be added.

When adding a unit to a target, under the hood a symbolic link is created in the target directory in /etc/systemd/system. If, for instance, you've enabled the vsftpd service to be automatically started, you'll find that a symbolic link /etc/systemd/system/multi-user.target/wants/vsftpd.service has been added, pointing to the unit file in /usr/lib/systemd/system/vsftpd.service and thus ensuring that the unit will automatically be started. In Systemd terminology, this symbolic link is known as a *want*, because it defines what the target wants to start when it is processed.

Managing Units Through Systemd

Managing the current state of Systemd units is an important task of a RHEL administrator. Managing units means not only managing their current state, but also changing options used by the different units.

As an administrator, managing Systemd units starts with starting and stopping units. You use the **systemctl** command to do that. In Exercise 11-1, you start, stop, and manage a unit. After you have configured a unit so that it can be started without problems, you need to make sure that it restarts automatically upon reboot. You do this by enabling or disabling the unit.

TIP The **systemctl** command has a large number of options, which may appear overwhelming at first sight, but there's no need to be overwhelmed. Just ensure that the **bash-completion** package is installed and use Tab completion on the **systemctl** command, which provides easy access to all of the available options.

Exercise 11-1 Managing Units with systemctl

1. Type **yum -y install vsftpd** to install the Very Secure FTP service.
2. Type **systemctl start vsftpd** to activate the FTP server on your machine.
3. Type **systemctl status vsftpd**. You'll get output as shown in Example 11-6, where you can see that the vsftpd service is currently operational. You can also see in the Loaded line that the service is currently disabled, which means that it will not be activated on a system restart. The vendor preset also shows as disabled, which means that by default after installation this unit will not automatically be enabled.
4. Type **systemctl enable vsftpd**. This creates a symbolic link in the wants directory for the multiuser target to ensure that the service is automatically started after a restart.
5. Type **systemctl status vsftpd** again. You'll now see that the unit file has changed from being disabled to enabled.

Key
Topic

Example 11-6 Requesting Current Unit Status with **systemctl status**

```
[root@server1 ~]# systemctl status vsftpd
vsftpd.service - Vsftpd ftp daemon
Loaded: loaded (/usr/lib/systemd/system/vsftpd.service; disabled;
       vendor preset: disabled)
Active: active (running) since Mon 2019-06-10 06:44:18 EDT; 5s ago
Process: 10044 ExecStart=/usr/sbin/vsftpd /etc/vsftpd/vsftpd.conf
       (code=exited, status=0/SUCCESS)
Main PID: 10045 (vsftpd)
Tasks: 1 (limit: 11365)
Memory: 696.0K
CGroup: /system.slice/vsftpd.service
        └─10045 /usr/sbin/vsftpd /etc/vsftpd/vsftpd.conf

Jun 10 06:44:18 server1.example.com systemd[1]: Starting Vsftpd ftp
daemon...
Jun 10 06:44:18 server1.example.com systemd[1]: Started Vsftpd ftp
daemon.
```

When requesting the current status of a Systemd unit as in Example 11-6, you can see different kinds of information about it. Table 11-2 shows the different kinds of information that you can get about unit files when using the **systemctl status** command.

Table 11-2 Systemd Status Overview

Status	Description
Loaded	The unit file has been processed and the unit is active.
Active(running)	The unit is running with one or more active processes.
Active(exited)	The unit has successfully completed a one-time run.
Active(waiting)	The unit is running and waiting for an event.
Inactive(dead)	The unit is not running.
Enabled	The unit will be started at boot time.
Disabled	The unit will not be started at boot time.
Static	The unit cannot be enabled but may be started by another unit automatically.

As an administrator, you’ll also often need to get a current overview of the status of Systemd unit files. Different commands can help you to get this insight, some of which are shown in Table 11-3.

Table 11-3 **systemctl** Unit Overview Commands

Command	Description
systemctl --type=service	Shows only service units
systemctl list-units --type=service	Shows all active service units (same result as the previous command)
systemctl list-units --type=service --all	Shows inactive service units as well as active service units
systemctl --failed --type=service	Shows all services that have failed
systemctl status -l your.service	Shows detailed status information about services

Managing Dependencies

In general, there are two ways to manage Systemd dependencies:

- Unit types such as socket and path are directly related to a service unit. Accessing either of these unit types will automatically trigger the service type.
- Dependencies can be defined within the unit, using keywords like **Requires**, **Requisite**, **After**, and **Before**.

As an administrator, you can request a list of unit dependencies. Type **systemctl list-dependencies** followed by a unit name to find out which dependencies it has, and add the **--reverse** option to find out which units are dependents of this unit. Example 11-7 shows an example of this command.

Example 11-7 Showing Unit Dependencies

```
[root@server1 ~]# systemctl list-dependencies vsftpd
vsftpd.service
└─system.slice
└─basic.target
   └─alsa-restore.service
   └─alsa-state.service
   └─firewalld.service
   └─microcode.service
   └─rhel-autorelabel-mark.service
   └─rhel-autorelabel.service
   └─rhel-configure.service
   └─rhel-dmesg.service
   └─rhel-loadmodules.service
```

```

└─paths.target
└─slices.target
  └─.slice
    └─system.slice
└─sockets.target
  └─avahi-daemon.socket
  └─cups.socket
  └─dbus.socket
  └─dm-event.socket
  └─iscsid.socket
  └─iscsiuio.socket
  └─lvm2-lvmetad.socket
  └─rpcbind.socket
  └─systemd-initctl.socket
  └─systemd-journald.socket
  └─systemd-shutdown.socket
  └─systemd-udev-control.socket
  └─systemd-udev-kernel.socket
└─sysinit.target
  └─dev-hugepages.mount
  └─dev-mqueue.mount
  └─dmraid-activation.service
  └─iscsi.service

```

To ensure accurate dependency management, different keywords can be used in the [Unit] section of a unit:

- **Requires:** If this unit loads, units listed here will load also. If one of the other units is deactivated, this unit will also be deactivated.
- **Requisite:** If the unit listed here is not already loaded, this unit will fail.
- **Wants:** This unit wants to load the units that are listed here, but it will not fail if any of the listed unit fails.
- **Before:** This unit will start before the unit specified with **Before**.
- **After:** This unit will start after the unit specified with **After**.

In upcoming Exercise 11-2 you'll learn how to use these options to manage unit dependency relations.

Managing Unit Options

When working with Systemd unit files, you risk getting overwhelmed with options. Every unit file can be configured with different options. To figure out which options are available for a specific unit, use the **systemctl show** command. For instance, the **systemctl show sshd** command shows all Systemd options that can be configured in the `sshd.service` unit, including their current default values. Example 11-8 shows the output of this command.

Example 11-8 Showing Available Options with **systemctl show**

```
[root@server1 ~]# systemctl show | head -20
Id=sshd.service
Names=sshd.service
Requires=basic.target
Wants=sshd-keygen.service system.slice
WantedBy=multi-user.target
ConsistsOf=sshd-keygen.service
Conflicts=shutdown.target
ConflictedBy=sshd.socket
Before=shutdown.target multi-user.target
After=network.target sshd-keygen.service systemd-journald.socket
    basic.target system.slice
Description=OpenSSH server daemon
LoadState=loaded
ActiveState=active
SubState=running
FragmentPath=/usr/lib/systemd/system/sshd.service
UnitFileState=enabled
InactiveExitTimestamp=Sat 2015-05-02 11:06:02 EDT
InactiveExitTimestampMonotonic=2596332166
ActiveEnterTimestamp=Sat 2015-05-02 11:06:02 EDT
ActiveEnterTimestampMonotonic=2596332166
ActiveExitTimestamp=Sat 2015-05-02 11:05:22 EDT
ActiveExitTimestampMonotonic=2559916100
InactiveEnterTimestamp=Sat 2015-05-02 11:06:02 EDT
```

When changing unit files to apply options, you need to make sure that the changes are written to `/etc/systemd/system`, which is the location where custom unit files should be created. The recommended way to do so is by using the **systemctl edit** command. This creates a subdirectory in `/etc/systemd/system` for the service that you are editing; for example, if you use **systemctl edit sshd.service**, you get a

directory with the name `/etc/systemd/systemd/ssh.service.d` in which a file with the name `override.conf` is created. All settings that are applied in this file overwrite any existing settings in the service file in `/usr/lib/systemd/system`. In Exercise 11-2 you learn how to apply changes to Systemd units.

An alternative (but definitely not recommended) way of making changes to unit files is to copy the default unit file from `/usr/lib/systemd/system` to `/etc/systemd/system`, and make your settings in the unit file in this location. At least this approach ensures that your modifications won't get lost after the default unit file gets updated.

TIP By default, Systemd uses the **nano** editor. Not everybody likes that very much (including me). If you want **vim** to be used instead of **nano**, edit the `/root/.profile` file to include the following line: **export SYSTEMD_EDITOR="/bin/vim"**. After logging in again, **vim** will be used as the default editor.

Exercise 11-2 Changing Unit Configuration

1. Type **yum install httpd** to install the Apache web server package.
2. Type **systemctl cat httpd.service** to show the current configuration of the unit file that starts the Apache web server.
3. Type **systemctl show httpd.service** to get an overview of available configuration options for this unit file.
4. Type **systemctl edit httpd.service** to change the default configuration, and ensure that the [Unit] section includes the lines **Restart=always** and **RestartSec=5s**.
5. Type **systemctl daemon-reload** to ensure that Systemd picks up the new configuration.
6. Type **systemctl restart httpd** to restart the httpd service.
7. Type **systemctl status httpd** and then repeat after 5 seconds. You'll notice that the httpd process gets automatically restarted.

Summary

In this chapter you learned how to work with Systemd. You've read how to manage Systemd service state and how to change different options in Systemd. In the next chapter you'll learn how to schedule tasks using the cron and at services.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 11-4 lists a reference of these key topics and the page numbers on which each is found.

Table 11-4 Key Topics for Chapter 11

Key Topic Element	Description	Page Number
Example 11-1	Unit types in Systemd	256
List	Three sections of a Systemd unit file	257
Section	Understanding Systemd Target Units	259
Exercise 11-1	Managing units with systemctl	261
Table 11-3	systemctl unit overview commands	263



Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

unit, wants, target, Systemd, dependencies

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. What is a unit?
2. Which command should you use to show all service units that are currently loaded?
3. How do you create a want for a service?
4. How do you change the default editor for **systemctl**?
5. Which directory contains custom Systemd unit files?
6. What should you include to ensure that a unit file will automatically load another unit file?
7. Which command will show available configuration options for the `httpd.service` unit?
8. Which command shows all dependencies for a specific unit?
9. What does it mean if **systemctl status** shows that a unit is dead?
10. How do you create a Systemd override file?

End-of-Chapter Lab

You have now learned how to work with Systemd. Before continuing, it is a good idea to work on a lab that helps you ensure that you can apply the skills that you acquired in this chapter.

Lab 11.1

1. Install the `vsftpd` and `httpd` services.
2. Set the default **systemctl** editor to **vim**.
3. Edit the `httpd.service` unit file such that starting `httpd` will always auto-start `vsftpd`. Edit the `httpd` service such that after failure it is automatically started again in 10 seconds.

This page intentionally left blank



The following topics are covered in this chapter:

- Configuring Cron to Automate Recurring Tasks
- Configuring At to Schedule Future Tasks

The following RHCSA exam objective is covered in this chapter:

- Schedule tasks using at and cron

On a Linux server it is important that certain tasks run at certain times. This can be done by using the At and Cron services, which can be configured to run tasks in the future. The At service is for executing future tasks once only, and the Cron service is for recurring regular tasks. In this chapter you learn how to configure both.

Scheduling Tasks

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 12-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 12-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Configuring cron to Automate Recurring Tasks	1–8
Configuring at to Schedule Future Tasks	9–10

- Which of the following commands enables you to check the current status of the crond service?
 - service crond status**
 - systemctl status crond**
 - /usr/sbin/crond --status**
 - chkconfig crond --show**
- Which of the following would run a cron task Sunday at 11 a.m.?
 - * 11 7 * ***
 - 0 11 * 7 ***
 - 0 11 * * 7**
 - 11 0 * 7 ***

3. Which of the following launches a job every five minutes from Monday through Friday?
 - a. `* /5 * * * 1-5`
 - b. `* /5 * 1-5 * *`
 - c. `0 /5 * * * 1-5`
 - d. `0 /5 * 1-5 * *`
4. How do you create a cron job for a specific user?
 - a. Log in as that user and type **`crontab -e`** to open the cron editor.
 - b. Open the crontab file in the user home directory and add what you want to add.
 - c. As root, type **`crontab -e username`**.
 - d. As root, type **`crontab -u username -e`**.
5. Which directory is mainly used by cron files that are installed automatically through RPM?
 - a. `/etc/crond.d`
 - b. `/etc/cron.d`
 - c. `/var/cron`
 - d. `/var/spool/cron`
6. Which of the following is not a recommended way to specify jobs that should be executed with cron?
 - a. Modify `/etc/crontab`.
 - b. Put the jobs in separate scripts in `/etc/cron.d`.
 - c. Use **`crontab -e`** to create user-specific cron jobs.
 - d. Put scripts in `/etc/cron.{hourly|daily|weekly|monthly}` for automatic execution.
7. Which service takes care of executing cron jobs in `/etc/cron.hourly`, `cron.daily`, `cron.weekly`, and `cron.monthly`?
 - a. `cron`
 - b. `crontab`
 - c. `atd`
 - d. `anacron`

8. Which of the statements about cron security is true?
 - a. By default, all users are allowed to schedule tasks through cron because the `/etc/cron.allow` file has the keyword **all** in it.
 - b. If the `cron.deny` file exists, a `cron.allow` file must be created also and list users who are allowed to schedule tasks through cron.
 - c. For every user, a matching entry must exist in either the `cron.allow` file or the `cron.deny` file.
 - d. If the `cron.allow` file exists, a user must be listed in it to be able to schedule jobs through cron.
9. After entering commands in the at shell, which command enables you to close the at shell?
 - a. Ctrl-V
 - b. Ctrl-D
 - c. **exit**
 - d. **:wq**
10. Which command enables you to see current at jobs scheduled for execution?
 - a. **atrm**
 - b. **atls**
 - c. **atq**
 - d. **at**

Foundation Topics

Configuring Cron to Automate Recurring Tasks

On a Linux system, some tasks have to be automated to occur on a regular basis. It would be one option to configure each process with a process-specific solution to handle recurring tasks, but that would not be efficient to deal with. That is why on Linux the cron service is used as a generic service to run processes automatically at specific times.

NOTE In RHEL 8, tasks can also be scheduled using a Systemd timer. However, because Systemd timers are not included in the RHCSA exam objectives, they are not covered in this chapter.

The cron service consists of two major components. First, there is the cron daemon **crond**. This daemon looks every minute to see whether there is work to do. Second, this work to do is defined in the cron configuration, which consists of multiple files working together to provide the right information to the right service at the right time. In this section, you learn how to configure cron.

Managing the **crond** Service

The **crond** service is started by default on every RHEL system. The service is needed because some system tasks run through **crond**. An example of these is **logrotate**, a service that cleans up log files and runs on a regular basis, but other important maintenance processes are started automatically through cron also.

Managing the **crond** service itself is easy: It does not need much management. Where other services need to be reloaded or restarted to activate changes to their configuration, this is not needed by **crond**. The **crond** daemon wakes up every minute and checks its configuration to see whether anything needs to be started.

To monitor the current status of the cron service, you can use the **systemctl status crond -l** command. Example 12-1 shows the output of this command.

Example 12-1 Monitoring the Current State of the crond Service

```
[root@server1 ~]# systemctl status crond -l
crond.service - Command Scheduler
Loaded: loaded (/usr/lib/systemd/system/crond.service; enabled;
       vendor preset: enabled)
Active: active (running) since Sat 2019-06-08 03:34:57 EDT;
       28min ago
Main PID: 1101 (crond)
Tasks: 2 (limit: 11363)
Memory: 2.3M
CGroup: /system.slice/crond.service
        └─1101 /usr/sbin/crond -n
          └─3600 /usr/sbin/anacron -s

Jun 08 03:34:57 server1.example.com crond[1101]: (CRON) STARTUP
(1.5.2)
Jun 08 03:34:57 server1.example.com crond[1101]: (CRON) INFO (Syslog
will be used instead of sendmail>
Jun 08 03:34:57 server1.example.com crond[1101]: (CRON) INFO (RANDOM_
DELAY will be scaled with factor>
Jun 08 03:34:58 server1.example.com crond[1101]: (CRON) INFO (running
with inotify support)
Jun 08 04:01:01 server1.example.com CROND[3591]: (root) CMD (run-parts
/etc/cron.hourly)
Jun 08 04:01:01 server1.example.com anacron[3600]: Anacron started on
2019-06-08
Jun 08 04:01:01 server1.example.com anacron[3600]: Will run job
'cron.daily' in 33 min.
Jun 08 04:01:01 server1.example.com anacron[3600]: Will run job
'cron.weekly' in 53 min.
Jun 08 04:01:01 server1.example.com anacron[3600]: Will run job
'cron.monthly' in 73 min.
Jun 08 04:01:01 server1.example.com anacron[3600]: Jobs will be
executed sequentially
```

The most significant part of the output of the **systemctl status crond** command is in the beginning, which indicates that the cron service is loaded and enabled. The fact that the service is enabled means that it will automatically be started whenever this service is restarting. The last part of the command output shows current status information. Through the journald service, the **systemctl** command can find out what is actually happening to the crond service.

Understanding Cron Timing

When scheduling services through Cron, you need to specify when exactly the services need to be started. In the crontab configuration (which is explained in more depth in the next section), you use a time string to indicate when tasks should be started. Table 12-2 shows the time and date fields used (in the order specified).

Key Topic

Table 12-2 cron Time and Date Fields

Field	Values
minute	0–59
hour	0–23
day of month	1–31
month	1–12 (or names that are better avoided)
day of week	0–7 (Sunday is 0 or 7, or names [which are better avoided])

In any of these fields, you can use an `*` as a wildcard to refer to any value. Ranges of numbers are allowed, as are lists and patterns. Some examples are listed next:

Key Topic

- `* 11 * * *` Every minute between 11:00 and 11:59 (probably not what you want)
- `0 11 * * 1-5` Every day at 11 a.m. on weekdays only
- `0 7-18 * * 1-5` Every hour at the top of the hour between 7 a.m. and 6 p.m. on weekdays
- `0 */2 2 12 5` Every 2 hours on the hour on December 2 and every Friday in December

TIP You don’t need to try to remember all this; `man 5 crontab` shows all possible constructions.

Managing Cron Configuration Files

The main configuration file for Cron is `/etc/crontab`, but you will not change this file directly. It does give you a convenient overview, though, of the different time specifications that can be used in Cron. It also sets environment variables that are used by the commands that are executed through Cron (see Example 12-2). To make modifications to the Cron jobs, there are other locations where cron jobs should be specified.

Example 12-2 /etc/crontab Sample Content

```
[root@server2 ~]# cat /etc/crontab
SHELL=/bin/bash
PATH=/sbin:/bin:/usr/sbin:/usr/bin
MAILTO=root

# For details see man 4 crontabs

# Example of job definition:
# .----- minute (0 - 59)
# | .----- hour (0 - 23)
# | | .----- day of month (1 - 31)
# | | | .----- month (1 - 12) OR jan,feb,mar,apr ...
# | | | | .---- day of week (0 - 6) (Sunday=0 or 7) OR
sun,mon,tue,wed,thu,fri,sat
# | | | | |
# * * * * * user-name command to be executed
```

Instead of modifying /etc/crontab, different Cron configuration files are used:

- Cron files in /etc/cron.d
- Scripts in /etc/cron.hourly, cron.daily, cron.weekly, and cron.monthly
- User-specific files that are created with **crontab -e**

**Key
Topic**

In this section, you get an overview of these locations.

NOTE If you want to experiment with how Cron works, you should allow for a sufficient amount of time for the job to be executed. The crond service reads its configuration every minute, after which new jobs can be scheduled for execution on the next minute. So, if you want to make sure your job is executed as fast as possible, allow for a safe margin of 3 minutes between the moment you save the cron configuration and the execution time.

To start, Cron jobs can be started for specific users. To create a user-specific cron job, type **crontab -e** after logging in as that user, or as root type **crontab -e -u username**.

When you are using **crontab -e**, the **vi** editor opens and creates a temporary file. After you edit the cron configuration, the temporary file is moved to its final location in the directory /var/spool/cron. In this directory, a file is created for each user. These files should never be edited directly! When the file is saved by **crontab -e**, it is activated automatically.

Whereas in the early days of RHEL the `/etc/crontab` file was modified directly, on RHEL 8 you do not do that anymore. If you want to add cron jobs that are not bound to a specific user account (and which for that reason by default will be executed as root if not specified otherwise), you add these to the `/etc/cron.d` directory. Just put a file in that directory (the exact name does not really matter) and make sure that it meets the syntax of a typical cron job. In Example 12-3, you can see an example of the `/etc/cron.d/unbound-anchor` cron configuration file (which was inserted to the `/etc/cron.d` directory upon installation of the unbound Domain Name System [DNS] server).

Example 12-3 Example cron Job in `/etc/cron.d`

```
[root@server1 cron.d]# cat unbound-anchor
# Look to see whether the DNSSEC Root key got rolled, if so check
trust and update

10 3 1 * * unbound /usr/sbin/unbound-anchor -a /var/lib/unbound/
root.anchor -c /etc/unbound/icannbundle.pem
```

This example file contains three elements. First there is the time indication, which has the command start at 3:10 a.m. on the first day of every month. Then, the configuration indicates that the command has to be started as the unbound user. The last part has the actual command that needs to be started with some arguments that are specific to the command and specify how this command should be used.

The last way to schedule Cron jobs is through the following directories:

- `/etc/cron.hourly`
- `/etc/cron.daily`
- `/etc/cron.weekly`
- `/etc/cron.monthly`

In these directories, you typically find scripts (not files that meet the crontab syntax requirements) that are put in there from RPM package files. When opening these scripts, notice that no information is included about the time when the command should be executed. That is because the exact time of execution does not really matter. The only thing that does matter is that the job is launched once an hour, day, week, or month.

Understanding the Purpose of anacron

To ensure regular execution of the job, Cron uses the anacron service. This service takes care of starting the hourly, daily, weekly, and monthly cron jobs, no matter

at which exact time. To determine how this should be done, anacron uses the `/etc/anacrontab` file. Example 12-4 shows the contents of the `/etc/anacrontab` file, which is used to specify how anacron jobs should be executed.

Example 12-4 anacrontab Configuration

```
[root@server1 spool]# cat /etc/anacrontab
# /etc/anacrontab: configuration file for anacron

# See anacron(8) and anacrontab(5) for details.

SHELL=/bin/sh
PATH=/sbin:/bin:/usr/sbin:/usr/bin
MAILTO=root
# the maximal random delay added to the base delay of the jobs
RANDOM_DELAY=45
# the jobs will be started during the following hours only
START_HOURS_RANGE=3-22

#period in days  delay in minutes  job-identifier  command
1          5   cron.daily          nice run-parts  /etc/cron.daily
7          25  cron.weekly        nice run-parts  /etc/cron.weekly
@monthly  45   cron.monthly       nice run-parts  /etc/cron.monthly
```

In `/etc/anacrontab`, the jobs to be executed are specified in lines that contain four fields, as shown in Example 12-4. The first field specifies the frequency of job execution, expressed in days. The second field specifies how long anacron waits before executing the job, which is followed by the third field that contains a job identifier. The fourth field specifies the command that should be executed.

TIP Although it's useful to know how anacron works, it typically is not a service that is configured directly. The need to configure services through anacron is taken away by the `/etc/cron.hourly`, `cron.daily`, `cron.weekly`, and `cron.monthly` files.

NOTE It is not easy to get an overview of the cron jobs actually scheduled for execution. There is no single command that would show all currently scheduled cron jobs. The **crontab -l** command does list cron jobs, but only for the current user account.

Managing Cron Security

By default, all users can enter Cron jobs. It is possible to limit which user is allowed to schedule Cron jobs by using the `/etc/cron.allow` and `/etc/cron.deny` configuration files. If the `cron.allow` file exists, a user must be listed in it to be allowed to use Cron. If the `/etc/cron.deny` file exists, a user must not be listed in it to be allowed to set up Cron jobs. Both files should not exist on the same system at the same time. Only root can use Cron if neither file exists.

In Exercise 12-1, you apply some of the cron basics and schedule cron jobs using different mechanisms.

Exercise 12-1 Running Scheduled Tasks Through cron

1. Open a root shell. Type `cat /etc/crontab` to get an impression of the contents of the `/etc/crontab` configuration file.
2. Type `crontab -e`. This opens an editor interface that by default uses `vi` as its editor. Add the following line:

```
0 2 * * 1-5 logger message from root
```
3. Use the `vi` command `:wq!` to close the editing session and write changes.
4. Type `cd /etc/cron.hourly`. In this directory, create a script file with the name **eachhour** that contains the following line:

```
logger This message is written at $(date)
```
5. Use `chmod +x eachhour` to make the script executable; if you fail to make it executable, it will not work.
6. Enter the directory `/etc/cron.d` and in this directory create a file with the name **eachhour**. Put the following contents in the file:

```
11 * * * * root logger This message is written from /etc/cron.d
```
7. Save the modifications to the configuration file and continue to the next section. (For optimal effect, perform step 8 after a couple of hours.)
8. After a couple of hours, type `grep written /var/log/messages` and read the messages to verify correct cron operations.

Configuring at to Schedule Future Tasks

Whereas cron is used to schedule jobs that need to be executed on a regular basis, the atd service is available for services that need to be executed only once. On RHEL 8, the atd service is available by default, so all that needs to be done is scheduling jobs.

To run a job through the atd service, you would use the **at** command, followed by the time the job needs to be executed. This can be a specific time, as in **at 14:00**, but it can also be a time indication like **at teatime** or **at noon**. After you type this, the at shell opens. From this shell, you can type several commands that will be executed at the specific time that is mentioned. After entering the commands, press Ctrl-D to quit the at shell.

After scheduling jobs with at, you can use the **atq** command (*q* for *queue*) to get an overview of all jobs currently scheduled. It is also possible to remove current at jobs. To do this, use the **atrm** command, optionally followed by the number of the at job that you want to remove. In Exercise 12-2, you learn how to work with at to schedule jobs for execution at a specific time.

TIP The **batch** command works like **at**, but it's a bit more sophisticated. When using **batch**, you can specify that a job is only started when system performance parameters allow. Typically, that is when system load is lower than 0.8. This value is a bit low on modern multi-CPU systems, which is why the load value can be specified manually when starting **atd**, using the **-l** command-line option. Use, for instance, **atd -l 3.0** to make sure that no batch job is started when system load is higher than 3.0.

Exercise 12-2 Scheduling Jobs with at

1. Type **systemctl status atd**. In the line that starts with Loaded:, this command should show you that the service is currently loaded and enabled, which means that it is ready to start receiving jobs.
2. Type **at 15:00** (or replace with any time near to the time at which you are working on this exercise).
3. Type **logger message from at**. Press Ctrl-D to close the at shell.
4. Type **atq** to verify that the job has indeed been scheduled.

Summary

In this chapter, you learned how to schedule jobs for future execution. You also learned how to configure cron to execute jobs repeatedly at a specific time. You learned that different methods exist to tell cron when a job should be executed. In addition, you learned about the anacron service, which is used to make sure that the jobs in the directories `/etc/cron.{hourly|daily|weekly|monthly}` are indeed executed, even if the system temporarily is not available. At the end of this chapter, you learned how to use the `atd` service to schedule tasks to be executed once.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 12-3 lists a reference of these key topics and the page number on which each is found.



Table 12-3 Key Topics for Chapter 12

Key Topic Element	Description	Page
Table 12-2	crontab time indicators	276
List	crontab time indicators examples	276
List	Methods to enter crontab information	277

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

cron, anacron, at

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. Where do you configure a cron job that needs to be executed once every 2 weeks?
2. How do you specify the execution time in a cron job that needs to be executed twice every month, on the 1st and the 15th of the month at 2 p.m.?
3. How do you specify cron execution time for a job that needs to run every 2 minutes every day?
4. How do you specify a job that needs to be executed on September 19 and every Thursday in September?
5. Which three valid day indicators can you use to specify that a cron job needs to be executed on Sunday?
6. Which command enables you to schedule a cron job for user lisa?
7. How do you specify that user boris is never allowed to schedule jobs through cron?
8. You need to make sure that a job is executed every day, even if the server at execution time is temporarily unavailable. How do you do this?
9. Which service must be running to schedule at jobs?
10. Which command enables you to find out whether any current at jobs are scheduled for execution?

End-of-Chapter Lab

In this end-of-chapter lab, you work on at jobs and on cron jobs.

Lab 12.1

1. Create a cron job that performs an update of all software on your computer every evening at 11 p.m.
2. Schedule your machine to be rebooted at 3 a.m. tomorrow morning.



The following topics are covered in this chapter:

- Understanding System Logging
- Configuring rsyslogd
- Rotating Log Files
- Working with journald

The following RHCSA exam objectives are covered in this chapter:

- Locate and interpret system log files and journals
- Preserve system journals

Analyzing log files is an important system administrator task. If anything goes wrong on a Linux system, the answer is often in the log files. On RHEL 8, two different log systems are used side by side, and it is important to know which information can be found where. This chapter teaches you all about it. You learn how to read log files, how to configure rsyslogd and journald, and how to set up your system for log rotation so that you can prevent your disks from being completely filled up by services that are logging too enthusiastically.

Configuring Logging

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 13-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 13-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Understanding System Logging	1–3
Configuring rsyslogd	4–7
Rotating Log Files	8
Working with journald	9–10

1. Which of the following statements about journald is *not* true?
 - a. journald logs kernel messages.
 - b. journald writes to the journal, which by default does not persist between boots.
 - c. journald is a replacement of rsyslogd.
 - d. To read files from the journal, the **journalctl** command is used.
2. Which log would you read to find messages related to authentication errors?
 - a. /var/log/messages
 - b. /var/log/lastlog
 - c. /var/log/audit/audit.log
 - d. /var/log/secure

3. Which log would you read to find information that relates to SELinux events?
 - a. `/var/log/messages`
 - b. `/var/log/lastlog`
 - c. `/var/log/audit/audit.log`
 - d. `/var/log/secure`
4. What is the name of the rsyslogd configuration file?
 - a. `/etc/rsyslog.conf`
 - b. `/etc/sysconfig/rsyslogd.conf`
 - c. `/etc/sysconfig/rsyslog.conf`
 - d. `/etc/rsyslog.d/rsyslogd.conf`
5. You need to change the startup behavior of the rsyslogd service. Which of the following describes the recommended approach to do so?
 - a. Include the startup parameter in the main rsyslog configuration file.
 - b. Create a snap-in file in the directory `/etc/rsyslog.d` and specify the required parameters in there.
 - c. Change the Systemd unit file in `/usr/lib/systemd/system` to include the required startup parameter.
 - d. Use the `SYSLOGD_OPTIONS` line in the file `/etc/sysconfig/rsyslog` and include the startup parameter here.
6. In the `rsyslog.conf` file, which of the following destinations refers to a specific rsyslogd module?
 - a. `-/var/log/maillog`
 - b. `/var/log/messages`
 - c. `:omusrmsg:*`
 - d. `*`
7. Which facility is the best solution if you want to configure the Apache web server to log messages through rsyslog?
 - a. `daemon`
 - b. `apache`
 - c. `syslog`
 - d. `local0-7`

8. You want to maximize the file size of a log file to 10 MB. Where do you configure this?
 - a. Create a file in `/etc/logrotate.d` and specify the maximal size in that file.
 - b. Put the maximal size in the logrotate cron job.
 - c. Configure the destination with the maximal size option.
 - d. This cannot be done.
9. Which directory is used to store the journald journal?
 - a. `/var/log/journal`
 - b. `/var/run/journal`
 - c. `/run/log`
 - d. `/run/log/journal`
10. What do you need to do to make the journald journal persistent?
 - a. Create the directory `/var/log/journal`, set appropriate permissions, and reboot your machine.
 - b. Open `/etc/sysconfig/journal` and set the `PERSISTENT` option to **yes**.
 - c. Open the `/etc/systemd/journald.conf` file and set the `PERSISTENT` option to **yes**.
 - d. Create the `/var/log/journal` file and set appropriate permissions.

Foundation Topics

Understanding System Logging

Most services used on a Linux server write information to log files. This information can be written to different destinations, and there are multiple solutions to find the relevant information in system logs. No less than three different approaches can be used by services to write log information:

- **Direct write:** Some services write logging information directly to the log files—even some important services such as the Apache web server and the Samba file server.
- **rsyslogd:** rsyslogd is the enhancement of syslogd, a service that takes care of managing centralized log files. Syslogd has been around for a long time.
- **journald:** With the introduction of Systemd, the journald log service systemd-journal has been introduced also. This service is tightly integrated with Systemd, which allows administrators to read detailed information from the journal while monitoring service status using the **systemctl status** command.

Understanding the Role of rsyslogd and journald

Key Topic

On RHEL 8, journald (which is implemented by the systemd-journal daemon) provides an advanced log management system. journald collects messages from the kernel, the entire boot procedure, and services and writes these messages to an event journal. This event journal is stored in a binary format, and it can be queried using the **journalctl** command. The **journalctl** command enables you to access a deep level of detail about messages that are logged, as it is an integrated part of Systemd and, as such, receives all messages that have been generated by Systemd units.

Key Topic

Because the journal that is written by journald is not persistent between reboots, messages are also forwarded to the rsyslogd service, which writes the messages to different files in the `/var/log` directory. rsyslogd also offers features that do not exist in journald, such as centralized logging and filtering messages by using modules. Numerous modules are available to enhance rsyslog logging, such as output modules that allow administrators to store messages in a database.

In the current state of RHEL 8, journald is not a replacement for rsyslog; it is just another way of logging information. journald is tightly integrated with Systemd; therefore, it logs everything that your server is doing. rsyslogd adds some services to it. In particular, it takes care of writing log information to specific files (that will be persistent between reboots), and it allows you to configure remote logging and log servers.

Apart from rsyslogd and systemd-journald, there is the auditd service. This service provides auditing, an in-depth trace of what specific services, processes, or users have been doing. Configuration of auditing is beyond the scope of the RHCSA exam, but you'll notice that SELinux, for instance, logs detailed messages to the auditd service.

To get more information about what has been happening on a machine running RHEL, administrators have to take three approaches:

- Monitor the files in /var/log that are written by rsyslogd.
- Use the **journalctl** command to get more detailed information from the journal.
- Use the **systemctl status <unit>** command to get a short overview of the last significant events that have been logged by Systemd units through journald. This command shows the status of services, as well as the last couple of lines that have been logged. Example 13-1 shows an example where this command clearly indicates what went wrong while starting a service.

Example 13-1 Using **systemctl status** to Show Relevant Log Information

```
[root@server1 ~]# systemctl status sshd -l
sshd.service - OpenSSH server daemon
Loaded: loaded (/usr/lib/systemd/system/sshd.service; enabled;
       vendor preset: enabled)
Active: active (running) since Sat 2019-06-08 03:34:56 EDT;
       55min ago
   Docs: man:sshd(8)
        man:sshd_config(5)
Main PID: 1055 (sshd)
   Tasks: 1 (limit: 11363)
  Memory: 5.5M
   CGroup: /system.slice/sshd.service
           └─1055 /usr/sbin/sshd -D -oCiphers=aes256-gcm@openssh.com,
             chacha20-poly1305@openssh.com,ae>

Jun 08 03:34:56 server1.example.com systemd[1]: Starting OpenSSH
server daemon...
Jun 08 03:34:56 server1.example.com sshd[1055]: Server listening on
0.0.0.0 port 22.
Jun 08 03:34:56 server1.example.com sshd[1055]: Server listening on ::
port 22.
Jun 08 03:34:56 server1.example.com systemd[1]: Started OpenSSH server
daemon.
Jun 08 03:57:38 server1.example.com sshd[3368]: Accepted password for
root from 192.168.4.1 port 5470>
Jun 08 03:57:38 server1.example.com sshd[3368]: pam_unix(sshd:
session):session opened for user root
```

Reading Log Files

Apart from the messages that are written by journald to the journal and which can be read using the **journalctl** command, on a Linux system you'll also find different log files in the directory `/var/log`. Most of the files in this directory are managed by rsyslogd, but some of the files are created directly by specific services. You can read these files by using a pager utility such as **less**.

The exact number of files in the `/var/log` directory will change, depending on the configuration of a server and the services that are running on that server. Some files, however, do exist on most occasions, and as an administrator, you should know which files they are and what kind of contents can be expected in these files.

Table 13-2 provides an overview of some of the standard files that are created in this directory.



Table 13-2 System Log Files Overview

Log File	Explanation
<code>/var/log/messages</code>	The most commonly used log file, it is the generic log file where most messages are written to.
<code>/var/log/dmesg</code>	Contains kernel log messages.
<code>/var/log/secure</code>	Contains authentication-related messages. Look here to see which authentication errors have occurred on a server.
<code>/var/log/boot.log</code>	Look here for messages that are related to system startup.
<code>/var/log/audit/audit.log</code>	Contains audit messages. SELinux writes to this file.
<code>/var/log/maillog</code>	Look here for mail-related messages.
<code>/var/log/samba</code>	Provides log files for the Samba service. Notice that Samba by default is not managed through rsyslog, but writes directly to the <code>/var/log</code> directory.
<code>/var/log/sss</code>	Contains messages that have been written by the sssd service, which plays an important role in the authentication process.
<code>/var/log/cups</code>	Contains log messages that were generated by the print service CUPS.
<code>/var/log/httpd/</code>	Directory that contains log files that are written by the Apache web server. Notice that Apache writes messages to these files directly and not through rsyslog.

Understanding Log File Contents

As an administrator, you need to be able to interpret the contents of log files. For example, Example 13-2 shows partial content from the `/var/log/messages` file.

Example 13-2 /var/log/messages Sample Content

```
[root@server1 ~]# tail -10 /var/log/messages
Jun  8 03:57:46 server1 journal[2682]: Received error from DBus
search provider org.gnome.Software.desktop: Gio.IOErrorEnum:
Timeout was reached
Jun  8 03:58:20 server1 journal[3078]: failed to get featured apps:
no apps to show
Jun  8 03:58:20 server1 journal[3078]: Only 0 apps for popular list,
hiding
Jun  8 03:58:20 server1 journal[2105]: Skipping refresh of rhel-8-for-
x86_64-highavailability-beta-source-rpms: cannot update repo 'rhel-
8-for-x86_64-highavailability-beta-source-rpms': Cannot download
repomd.xml: Curl error (42): Operation was aborted by an application
callback for https://cdn.redhat.com/content/beta/rhel8/8/x86_64/
highavailability/source/SRPMS/repodata/repomd.xml [Callback
aborted]; Last error: Curl error (42): Operation was aborted by
an application callback for https://cdn.redhat.com/content/beta/
rhel8/8/x86_64/highavailability/source/SRPMS/repodata/repomd.xml
[Callback aborted]
Jun  8 03:58:20 server1 journal[3078]: hiding category graphics
featured applications: found only 0 to show, need at least 9
Jun  8 03:58:23 server1 journal[2105]: Skipping refresh of rhel-8-
for-x86_64-supplementary-beta-debug-rpms: cannot update repo 'rhel-
8-for-x86_64-supplementary-beta-debug-rpms': Cannot download repomd.
xml: Curl error (42): Operation was aborted by an application
callback for https://cdn.redhat.com/content/beta/rhel8/8/x86_64/
supplementary/debug/repodata/repomd.xml [Callback aborted]; Last
error: Curl error (42): Operation was aborted by an application
callback for https://cdn.redhat.com/content/beta/rhel8/8/x86_64/
supplementary/debug/repodata/repomd.xml [Callback aborted]
Jun  8 03:58:46 server1 chronyd[907]: Selected source 5.200.6.34
Jun  8 03:59:16 server1 systemd[2571]: Starting Mark boot as
successful...
Jun  8 03:59:16 server1 systemd[2571]: Started Mark boot as
successful.
Jun  8 04:00:16 server1 systemd[3374]: Starting Mark boot as
successful...
```

As you can see in Example 13-2, each line that is logged has specific elements:

- **Date and time:** Every log message starts with a timestamp. For filtering purposes, the timestamp is written as military time.
- **Host:** The host the message originated from. This is relevant because rsyslogd can be configured to handle remote logging as well.
- **Service or process name:** The name of the service or process that generated the message.
- **Message content:** The content of the message, which contains the exact message that has been logged.

To read the content of a log file, you can use a pager utility, like **less**, or you can live monitor what is happening in the log file, as described in the next section.

Live Log File Monitoring

When you are configuring services on Linux, it might be useful to see in real time what is happening. You could, for example, open two terminal sessions at the same time. In one terminal session, you configure and test the service. In the other terminal session, you see in real time what is happening. The **tail -f <logfile>** command shows in real time which lines are added to the log file. Exercise 13-1 in the following section shows a small example in which **tail -f** is used. When monitoring a log file with **tail -f**, the trace remains open until you press Ctrl-C to close it.

Using logger

Most services write information to the log files all by themselves or through rsyslogd. The **logger** command enables users to write messages to rsyslog from the command line or a script. Using this command is simple. Just type **logger**, followed by the message you want to write to the logs. The **logger** utility, in this way, offers a convenient solution to write messages from scripts. This allows you to have a script write to syslog if something goes wrong.

When using **logger**, you can also specify the priority and facility to log to. The command **logger -p kern.err hello** writes **hello** to the kernel facility, for example, using the error priority (discussed later in this chapter). This option enables you to test the working of specific rsyslog facilities. In Exercise 13-1, you use **tail -f** to monitor a log file in real time and use **logger** to write log messages to a log file.

Exercise 13-1 Using Live Log Monitoring and logger

1. Open a root shell.
2. From the root shell, type **tail -f /var/log/messages**.
3. Open a second terminal window. In this terminal window, type **su - student** to open a subshell as user student.
4. Type **su -** to open a root shell, but enter the wrong password.
5. Notice that nothing appears in **/var/log/messages**. That is because login-related errors are not written here.
6. From the student shell, type **logger hello**. You'll see the message appearing in the **/var/log/messages** file in real time.
7. In the **tail -f** terminal, press Ctrl-C to stop tracing the messages file.
8. Type **tail -20 /var/log/secure**. This shows the last 20 lines in **/var/log/secure**, which also shows the messages that the **su -** password errors have generated previously.

Configuring rsyslogd

To make sure that the information that needs to be logged is written to the location where you want to find it, you can configure the rsyslogd service through the `/etc/rsyslog.conf` file. In this file, you find different sections that allow you to specify where and how information should be written.

Understanding rsyslogd Configuration Files

Like many other services on RHEL 8, the configuration for rsyslogd is not defined in just one configuration file. The `/etc/rsyslog.conf` file is the central location where rsyslogd is configured. From this file, the content of the directory `/etc/rsyslog.d` is included. This directory can be populated by installing RPM packages on a server. When looking for specific log configuration, make sure to always consider the contents of this directory also.

If specific options need to be passed to the rsyslogd service on startup, you can do this by using the `/etc/sysconfig/rsyslog` file. This file by default contains one line, which reads `SYSLOGD_OPTIONS=""`. On this line, you can specify rsyslogd startup parameters. The `SYSLOGD_OPTIONS` variable is included in the `Systemd` configuration file that starts rsyslogd. Theoretically, you could change startup parameters in this file, as well, but that is not recommended. (See Chapter 17, “Managing and Understanding the Boot Procedure,” for more details about `Systemd` configuration.)

Understanding rsyslog.conf Sections

The `rsyslog.conf` file is used to specify what should be logged and where it should be logged. To do this, you’ll find different sections in the `rsyslog.conf` file:

- **#### MODULES ####:** rsyslogd is modular. Modules are included to enhance the supported features in rsyslogd.
- **#### GLOBAL DIRECTIVES ####:** This section is used to specify global parameters, such as the location where auxiliary files are written or the default timestamp format.
- **#### RULES ####:** This is the most important part of the `rsyslog.conf` file. It contains the rules that specify what information should be logged to which destination.

Understanding Facilities, Priorities, and Log Destinations

To specify what information should be logged to which destination, rsyslogd uses facilities, priorities, and destinations:

- A *facility* specifies a category of information that is logged. rsyslogd uses a fixed list of facilities, which cannot be extended. This is because of backward compatibility with the legacy syslog service.
- A *priority* is used to define the severity of the message that needs to be logged. When specifying a priority, by default all messages with that priority and all higher priorities are logged.
- A *destination* defines where the message should be written. Typical destinations are files, but rsyslog modules can be used as a destination as well, to allow further processing through a rsyslogd module.

Example 13-3 shows an example of the RULES section in rsyslog.

Example 13-3 Example of the RULES Section in rsyslog.conf

```
#### RULES ####

# Log all kernel messages to the console.
# Logging much else clutters up the screen.
#kern.*                                     /dev/console

# Log anything (except mail) of level info or higher.
# Do not log private authentication messages!
*.info;mail.none;authpriv.none;cron.none  /var/log/messages

# The authpriv file has restricted access.
authpriv.*                                 /var/log/secure

# Log all the mail messages in one place.
mail.*                                     -/var/log/maillog

# Log cron stuff
cron.*                                     /var/log/cron

# Everybody gets emergency messages
*.emerg                                   :omusrmsg:*

# Save news errors of level crit and higher in a special file.
uucp,news.crit                            /var/log/spooler
```

In Example 13-3, you can see how different facilities and priorities are used to define locations where information can be logged. The available facilities and priorities are fixed and cannot be added to. Table 13-3 shows which facilities are available, and Table 13-4 shows a list of all priorities.

When specifying a destination, a file is often used. If the filename starts with a hyphen (as in `-/var/log/maillog`), the log messages will not be immediately committed to the file but instead will be buffered to make writes more efficient. Device files can also be used, such as `/dev/console`. If this device is used, messages are written in real time to the console. On modern servers, this often does not make sense, because administrators often log in remotely and do not see what is happening on the server console.

Table 13-3 rsyslogd Facilities

Facility	Used by
auth/authpriv	Messages related to authentication.
cron	Messages generated by the crond service.
daemon	Generic facility that can be used for nonspecified daemons.
kern	Kernel messages.
lpr	Messages generated through the legacy lpd print system.
mail	Email-related messages.
mark	Special facility that can be used to write a marker periodically.
news	Messages generated by the NNTP news system.
security	Same as auth/authpriv. Should not be used anymore.
syslog	Messages generated by the syslog system.
user	Messages generated in user space.
uucp	Messages generated by the legacy UUCP system.
local0-7	Messages generated by services that are configured by any of the local0 through local7 facilities.

**Key
Topic**

The syslog facilities were defined in the 1980s, and to guarantee backward compatibility, no new facilities can be added. The result is that some facilities still exist that basically serve no purpose anymore, and some services that have become relevant at a later stage do not have their own facility. As a solution, two specific facility types can be used. The daemon facility is a generic facility that can be used by any daemon. In addition, the local0 through local7 facilities can be used.

If there are services that do not have their own rsyslogd facility that need to write log messages to a specific log file anyway, these services can be configured to use any

of the local0 through local7 facilities. You next have to configure the services to use these facilities as well. The procedure you follow to do that is specific to the service you are using. Then you need to add a rule to the rsyslog.conf file to send messages that come in through that facility to a specific log file.

To determine which types of messages should be logged, different severities can be used in rsyslog.conf lines. These severities are the syslog priorities. Table 13-4 provides an overview of the available priorities in ascending order.

Key
Topic

Table 13-4 rsyslogd Priorities

Priority	Description
debug	Debug messages that will give as much information as possible about service operation.
info	Informational messages about normal service operation.
notice	Informational messages about items that might become an issue later.
warning / warn	Something is suboptimal, but there is no real error yet.
err /error	A noncritical error has occurred.
crit	A critical error has occurred.
alert	Used when the availability of the service is about to be discontinued.
emerg/panic	Message generated when the availability of the service is discontinued.

When a specific priority is used, all messages with that priority and higher are logged according to the specifications used in that specific rule. If you need to configure logging in a detailed way, where messages with different priorities are sent to different files, you can specify the priority with an equals sign (=) in front of it, as in the following line, which will write all cron messages with only the debug priority to a specific file with the name /var/log/cron.debug. The - in front of the line specifies to buffer writes so that information is logged in a more efficient way.

```
cron.=debug -/var/log/cron.debug
```

TIP There is no need to learn the names of rsyslogd facilities and priorities by heart. They are all listed in man 5 rsyslog.conf. On the exam, you have access to the man pages, so this information will be easily accessible.

Exercise 13-2 shows how to change rsyslog.conf. You configure the Apache service to log messages through syslog, and you create a rule that logs debug messages to a specific log file.

Exercise 13-2 Changing rsyslog.conf Rules

1. By default, the Apache service does not log through rsyslog but keeps its own logging. You are going to change that. To start, type **yum install -y httpd** to install the Apache service.
2. After installing the Apache service, open its configuration file `/etc/http/conf/httpd.conf` and add the following line to it:
ErrorLog syslog:local1
3. Type **systemctl restart httpd**.
4. Create a line in the `rsyslog.conf` file that will send all messages that it receives for facility `local1` (which is now used by the `httpd` service) to the file `/var/log/httpd-error.log`. To do this, include the following line:
local1.error /var/log/httpd-error.log
5. Tell `rsyslogd` to reload its configuration, by using **systemctl restart rsyslog**.
6. All Apache error messages will now be written to the `httpd-error.log` file.
7. From the Firefox browser, go to `http://localhost`. As no `index.html` page exists yet, this will be written to the error log.
8. Create a snap-in file that logs debug messages to a specific file as well. To do this, type **echo "*.debug /var/log/messages-debug">/etc/rsyslog.d/debug.conf**.
9. Again, restart `rsyslogd` using **systemctl restart rsyslog**.
10. Use the command **tail -f /var/log/messages-debug** to open a trace on the newly created file.
11. Type **logger -p daemon.debug "Daemon Debug Message"**. You'll see the debug message passing by.
12. Press **Ctrl-C** to close the debug log file.

Rotating Log Files

To prevent `syslog` messages from filling up your system completely, the log messages can be rotated. That means that when a certain threshold has been reached, the old log file is closed and a new log file is opened. The `logrotate` utility is started periodically through the `crond` service to take care of rotating log files.

When a log file is rotated, the old log file is typically copied to a file that has the rotation date in it. So, if `/var/log/messages` is rotated on June 8, 2019, the rotated filename will be `/var/log/messages-20190608`. As a default, four old log files are kept on the system. Files older than that period are removed from the system automatically.

WARNING Log files that have been rotated are not stored anywhere; they are just gone. If your company policy requires you to be able to access information about events that have happened more than five weeks ago, for example, you should either back up log files or configure a centralized log server where logrotate keeps rotated messages for a significantly longer period.

The default settings for log rotation are kept in the file `/etc/logrotate.conf` (see Example 13-4).

Example 13-4 `/etc/logrotate.conf` Sample Content

```
# see "man logrotate" for details
# rotate log files weekly
weekly

# keep 4 weeks worth of backlogs
rotate 4

# create new (empty) log files after rotating old ones
create

# use date as a suffix of the rotated file
dateext

# uncomment this if you want your log files compressed
#compress

# RPM packages drop log rotation information into this directory
include /etc/logrotate.d

# system-specific logs may be also be configured here
```

The most significant settings used in this configuration file tell logrotate to rotate files on a weekly basis and keep four old versions of the file. You can obtain more information about other parameters in this file through the **man logrotate** command.

If specific files need specific settings, you can create a configuration file for that file in `/etc/logrotate.d`. The settings for that specific file overwrite the default settings in `/etc/logrotate.conf`. You will find that different files exist in this directory already to take care of some of the configuration files.

Working with journald

The `systemd-journald` service stores log messages in the journal, a binary file that is temporarily stored in the file `/run/log/journal`. This file can be examined using the **journalctl** command.

Using journalctl to Find Events

The easiest way to use **journalctl** is by just typing the command. It shows that recent events have been written to the journal since your server last started. The result of this command is shown in the **less** pager, and by default you'll see the beginning of the journal. Because the journal is written from the moment your server boots, this shows boot-related log messages. If you want to see the last messages that have been logged, you can use **journalctl -f**, which shows the last lines of the messages where new log lines are automatically added. You can also type **journalctl** and use (uppercase) **G** to go to the end of the journal. Also note that the search options `/` and `?` work in the **journalctl** output. Example 13-5 shows a partial result of this command.

Example 13-5 Watching Log Information Generated by journald

```
-- Logs begin at Sat 2019-06-08 04:45:34 EDT, end at Sat 2019-06-08
04:56:11 EDT. --
Jun 08 04:45:34 server1.example.com kernel: Linux version 4.18.0-80.
e18.x86_64 (mockbuild@x86-vm-08.b>
Jun 08 04:45:34 server1.example.com kernel: Command line: BOOT_
IMAGE=(hd0,msdos1)/vmlinuz-4.18.0-80.e>
Jun 08 04:45:34 server1.example.com kernel: Disabled fast string
operations
Jun 08 04:45:34 server1.example.com kernel: x86/fpu: Supporting XSAVE
feature 0x001: 'x87 floating po>
Jun 08 04:45:34 server1.example.com kernel: x86/fpu: Supporting XSAVE
feature 0x002: 'SSE registers'
```

```

Jun 08 04:45:34 server1.example.com kernel: x86/fpu: Supporting XSAVE
feature 0x004: 'AVX registers'
Jun 08 04:45:34 server1.example.com kernel: x86/fpu: Supporting XSAVE
feature 0x008: 'MPX bounds regi>
Jun 08 04:45:34 server1.example.com kernel: x86/fpu: Supporting XSAVE
feature 0x010: 'MPX CSR'
Jun 08 04:45:34 server1.example.com kernel: x86/fpu: xstate_offset[2]:
576, xstate_sizes[2]: 256
Jun 08 04:45:34 server1.example.com kernel: x86/fpu: xstate_offset[3]:
832, xstate_sizes[3]: 64
Jun 08 04:45:34 server1.example.com kernel: x86/fpu: xstate_offset[4]:
896, xstate_sizes[4]: 64
Jun 08 04:45:34 server1.example.com kernel: x86/fpu: Enabled xstate
features 0x1f, context size is 96>

```

What makes **journalctl** a flexible command is that its many filtering options allow you to show exactly what you need. Exercise 13-3 shows some of the most interesting options.

Exercise 13-3 Discovering journalctl

1. Type **journalctl**. You'll see the content of the journal since your server last started, starting at the beginning of the journal. The content is shown in **less**, so you can use common **less** commands to walk through the file.
2. Type **q** to quit the pager. Now type **journalctl --no-pager**. This shows the contents of the journal without using a pager.
3. Type **journalctl -f**. This opens the live view mode of **journalctl**, which allows you to see new messages scrolling by in real time. Press Ctrl-C to interrupt.
4. Type **journalctl** and press the Tab key twice. This shows specific options that can be used for filtering. Type, for instance, **journalctl _UID=1000** to show messages that have been logged for your student user account.
5. Type **journalctl -n 20**. The **-n 20** option displays the last 20 lines of the journal (just like **tail -n 20**).
6. Type **journalctl -p err**. This command shows errors only.
7. If you want to view journal messages that have been written in a specific time period, you can use the **--since** and **--until** commands. Both options take the time parameter in the format YYYY-MM-DD hh:mm:ss. Also, you can use **yesterday**, **today**, and **tomorrow** as parameters. So, type **journalctl --since yesterday** to show all messages that have been written since yesterday.

8. **journalctl** allows you to combine different options, as well. So, if you want to show all messages with a priority **err** that have been written since yesterday, use **journalctl --since yesterday -p err**.
9. If you need as much detail as possible, use **journalctl -o verbose**. This shows different options that are used when writing to the journal (see Example 13-6). All these options can be used to tell the **journalctl** command which specific information you are looking for. Type, for instance, **journalctl _SYSTEMD_UNIT=sshd.service** to show more information about the **sshd** Systemd unit.
10. Type **journalctl --dmesg**. This shows kernel-related messages only. Not many people use this command, as the **dmesg** command gives the exact same result.

In the preceding exercise, you typed **journalctl -o verbose** to show verbose output. Example 13-6 shows an example of the **verbose** output. As you can see, this provides detailed information for all items that have been logged, including the PID, the ID of the associated user and group account, the command that is associated, and more. This verbose information may help you in debugging specific Systemd units.

Example 13-6 Showing Detailed Log Information with **journalctl -o verbose**

```
[root@server1 ~]# journalctl _SYSTEMD_UNIT=sshd.service -o verbose
-- Logs begin at Sat 2019-06-08 04:45:34 EDT, end at Sat 2019-06-08
05:01:40 EDT. --
Sat 2019-06-08 04:45:52.633752 EDT [s=53e57e2481434e078e8306367dc5645c;
i=898;b=f35bb68348284f9ead79c3>
  _BOOT_ID=f35bb68348284f9ead79c3c6750adfa1
  _MACHINE_ID=5aa095b495ed458d934c54a88078c165
  _HOSTNAME=server1.example.com
  PRIORITY=6
  _UID=0
  _GID=0
  _SYSTEMD_SLICE=system.slice
  _CAP_EFFECTIVE=3fffffffff
  _TRANSPORT=syslog
  SYSLOG_FACILITY=10
  SYSLOG_IDENTIFIER=sshd
  SYSLOG_PID=1211
  MESSAGE=Server listening on 0.0.0.0 port 22.
  _PID=1211
  _COMM=sshd
```



```

_EXE=/usr/sbin/sshd
_CMDLINE=/usr/sbin/sshd -D -oCiphers=aes256-gcm@openssh.com,
chacha20-poly1305@openssh.com,aes256->
_SELINUX_CONTEXT=system_u:system_r:sshd_t:s0-s0:c0.c1023
_SYSTEMD_CGROUP=/system.slice/sshd.service
_SYSTEMD_UNIT=sshd.service
_SYSTEMD_INVOCATION_ID=728a7dfecd7d436387dcd6e319c208c7
_SOURCE_REALTIME_TIMESTAMP=1559983552633752
Sat 2019-06-08 04:45:52.634696 EDT [s=53e57e2481434e078e8306367dc5645c;
i=899;b=f35bb68348284f9ead79c3>
_BOOT_ID=f35bb68348284f9ead79c3c6750adfa1
lines 1-26

```

Preserving the Systemd Journal

By default, the journal is stored in the file `/run/log/journal`. The entire `/run` directory is used for current process status information only, which means that the journal is cleared when the system reboots. To make the journal persistent between system restarts, you should make sure that a directory `/var/log/journal` exists.

Storing the journal permanently requires setting the `Storage=auto` parameter in `/etc/systemd/journal.conf`. This parameter can have different values:

- **Storage=auto** The journal will be written on disk if the directory `/var/log/journal` exists.
- **Storage=volatile** The journal will be stored only in the `/run/log/` journal directory.
- **Storage=persistent** The journal will be stored on disk in the directory `/var/log/journal`. This directory will be created automatically if it doesn't exist.
- **Storage=none** No data will be stored, but forwarding to other targets such as the kernel log buffer or syslog will still work.

Even when the journal is written to the permanent file in `/var/log/journal`, that does not mean that the journal is kept forever. The journal has built-in log rotation that will be used monthly. Also, the journal is limited to a maximum size of 10% of the size of the file system that it is on, and it will stop growing if less than 15% of the file system is still free. If that happens, the oldest messages from the journal are dropped automatically to make room for newer messages. To change these settings, you can modify the file `/etc/systemd/journald.conf`, as shown in Example 13-7 (along with other parameters you can set).

Example 13-7 Setting journald Parameters Through /etc/systemd/journal.conf

```
[Journal]
#Storage=auto
#Compress=yes
#Seal=yes
#SplitMode=uid
#SyncIntervalSec=5m
#RateLimitIntervalSec=30s
#RateLimitBurst=10000
#SystemMaxUse=
#SystemKeepFree=
#SystemMaxFileSize=
#SystemMaxFiles=100
#RuntimeMaxUse=
#RuntimeKeepFree=
#RuntimeMaxFileSize=
#RuntimeMaxFiles=100
#MaxRetentionSec=
#MaxFileSec=1month
#ForwardToSyslog=no
#ForwardToKMsg=no
#ForwardToConsole=no
#ForwardToWall=yes
#TTYPath=/dev/console
#MaxLevelStore=debug
#MaxLevelSyslog=debug
```

Making the journald journal permanent is not hard to do. Exercise 13-4 shows how to proceed.

Exercise 13-4 Making the journald Journal Permanent

1. Open a root shell and type **mkdir /var/log/journal**.
2. Before journald can write the journal to this directory, you have to set ownership. Type **chown root:systemd-journal /var/log/journal**, followed by **chmod 2755 /var/log/journal**.
3. Next, you can either reboot your system (restarting the systemd-journald service is not enough) or use the **killall -USR1 systemd-journald** command.
4. The Systemd journal is now persistent across reboots. If you want to see the log messages since the last reboot, use **journalctl -b**.

**Key
Topic**

Summary

In this chapter, you learned how to configure logging. You read how the `rsyslogd` and `journald` services are used on RHEL 8 to keep log information, and you learned how to manage logs that are written by these services. You also learned how to configure log rotation and make the journal persistent.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 13-5 lists a reference of these key topics and the page number on which each is found.

**Key
Topic**

Table 13-5 Key Topics for Chapter 13

Key Topic Element	Description	Page
Paragraph	<code>journald</code> explanation	288
Paragraph	<code>rsyslogd</code> explanation	288
Table 13-2	System log files overview	290
Table 13-3	<code>rsyslogd</code> facilities	295
Table 13-4	<code>rsyslogd</code> priorities	296
Exercise 13-4	Making the <code>journald</code> journal permanent	303

Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

journal, **journalctl**, rsyslogd, facility, priority, destination, log rotation

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. Which file is used to configure rsyslogd?
2. Which configuration file contains messages related to authentication?
3. If you do not configure anything, how long will it take for log files to be rotated away?
4. Which command enables you to log a message from the command line to the user facility, using the notice priority?
5. Which line would you add to write all messages with a priority of info to the file `/var/log/messages.info`?
6. Which configuration file enables you to allow the journal to grow beyond its default size restrictions?
7. Which command enables you to see new messages in the journal scrolling by in real time?
8. Which command enables you to see all journald messages that have been written for PID 1 between 9:00 a.m. and 3:00 p.m.?
9. Which command enables you to see journald messages since the last reboot on a system where a persistent journal has been configured?
10. Which procedure enables you to make the journald journal persistent?

End-of-Chapter Lab

You have now learned how to work with logging on Red Hat Enterprise Linux 8 and know how to configure rsyslogd and journald. You can now complete the end-of-chapter lab to reinforce these newly acquired skills.

Lab 13.1

1. Configure the journal to be persistent across system reboots.
2. Make a configuration file that writes all messages with an info priority to the file `/var/log/messages.info`.
3. Configure logrotate to keep ten old versions of log files.

This page intentionally left blank



The following topics are covered in this chapter:

- Understanding MBR and GPT Partitions
- Managing Partitions and File Systems
- Mounting File Systems

The following RHCSA exam objectives are covered in this chapter:

- List, create, delete partitions on MBR and GPT disks
- Configure systems to mount file systems at boot by universally unique ID (UUID) or label
- Add new partitions and logical volumes, and swap to a system nondestructively
- Create, mount, unmount, and use vfat, ext4, and xfs file systems

Working with storage is an important task for a Linux administrator. In this chapter, you acquire the first set of essential storage skills. You learn how to create and manage partitions, format them with the file system you need to use, and mount these file systems.

Managing Storage

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 14-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 14-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Understanding MBR and GPT Partitions	1–2
Managing Partitions and File Systems	3–6
Mounting File Systems	7–10

- Which of the following is *not* an advantage of using a GUID partition table over using an MBR partition table?
 - Access time to a directory is quicker.
 - A total amount of 8 ZiB can be addressed by a partition.
 - With GUID partitions, a backup copy of the partition table is created automatically.
 - There can be up to 128 partitions in total.
- You want to create a partition with a size of 1024^5 bytes. What size should it be?
 - 1 PB
 - 1 PiB
 - 1 EB
 - 1 EiB

3. Which partition type is commonly used to create a Linux partition?
 - a. 81
 - b. 82
 - c. 83
 - d. 8e
4. What is the default disk device name you would expect to see in KVM virtual machines?
 - a. /dev/sda
 - b. /dev/hda
 - c. /dev/sda
 - d. /dev/xsda
5. Which of the following statements is *not* true?
 - a. Do not ever use **gdisk** on an MBR disk.
 - b. **fdisk** offers support to manage GPT partitions as well.
 - c. Depending on your needs, you can create MBR and GPT partitions on the same disk.
 - d. If your server boots from EFI, you must use GPT partitions.
6. Which of the following file systems is used as the default in RHEL8?
 - a. Ext4
 - b. XFS
 - c. btrfs
 - d. Ext3
7. Which command enables you to find current UUIDs set to the file systems on your server?
 - a. **mount**
 - b. **df -h**
 - c. **lsblk**
 - d. **blkid**

8. What would you put in the device column of `/etc/fstab` to mount a file system based on its unique ID `42f419c4-633f-4ed7-b161-519a4dadd3da`?
 - a. `42f419c4-633f-4ed7-b161-519a4dadd3da`
 - b. `/dev/42f419c4-633f-4ed7-b161-519a4dadd3da`
 - c. `ID=42f419c4-633f-4ed7-b161-519a4dadd3da`
 - d. `UUID=42f419c4-633f-4ed7-b161-519a4dadd3da`
9. Which of the following `/etc/fstab` lines would perform a file system check on the file system, but only after the root file system has been checked successfully?
 - a. `/dev/sda1 /data xfs defaults 1 1`
 - b. `/dev/sda1 /data xfs defaults 1 2`
 - c. `/dev/sda1 /data xfs defaults 1 1`
 - d. `/dev/sda1 /data xfs defaults 0 2`
10. Which mount option would you use in `/etc/fstab` to specify that the file system can be mounted only after the network is available?
 - a. `network`
 - b. `_netdev`
 - c. `_network`
 - d. `netdev`

Foundation Topics

Understanding MBR and GPT Partitions

To use a hard drive, it needs to have partitions. Some operating systems install everything to one partition, while other operating systems such as Linux normally have several partitions on one hard disk. Using more than one partition on a system makes sense for multiple reasons:

- It's easier to distinguish between different types of data.
- Specific mount options can be used to enhance security or performance.
- It's easier to create a backup strategy where only relevant portions of the OS are backed up.
- If one partition accidentally fills up completely, the other partitions still are usable and your system might not crash immediately.

NOTE Instead of using multiple different partitions, you can also use different LVM logical volumes. Managing logical volumes is covered in Chapter 15, “Managing Advanced Storage.”

On RHEL 8, two different partitioning schemes are available. Before creating your first partition, you should understand these schemes and be able to determine which scheme works best in a specific environment.

Understanding the MBR Partitioning Scheme

When the personal computer was invented in the early 1980s, a system was needed to define hard disk layout. This system became known as the Master Boot Record (MBR) partitioning scheme. While booting a computer, the Basic Input Output System (BIOS) was loaded to access hardware devices. From the BIOS, the bootable disk device was read, and on this bootable device, the MBR was allocated. The MBR contains all that is needed to start a computer, including a boot loader and a partition table.

When hard disks first came out for PCs in the early 1980s, users could have different operating systems on them. Some of these included MS-DOS/PC-DOS, PC/IX (IBM's UNIX for 8086 PCs), CPM86, and MPM86. The disk would be partitioned so each operating system installed got a part of the disk. One of the partitions would be made active, meaning the code in the boot sector in the MBR would read the first sector of that active partition and run the code. That code would then load the rest of the OS. This explains why four partitions were deemed

“enough.” Some operating systems (such as SCO Xenix and SCO Unix) would have another layer of partitions in the dedicated UNIX partition (in SCO’s case, called *divisions*), with its own partition program (SCO: *divvy*)

The MBR was defined as the first 512 bytes on a computer hard drive, and in the MBR an operating system boot loader (such as GRUB 2; see Chapter 17, “Managing and Understanding the Boot Procedure”) was present, as well as a partition table. The size that was used for the partition table was relatively small, just 64 bytes, with the result that in the MBR no more than four partitions could be created. Since partition size data was stored in 32-bit values, and a default sector size of 512 bytes was used, the maximum size that could be used by a partition was limited to 2 TiB (hardly a problem in the early 1980s).

In the MBR, just four partitions could be created. Because many PC operating systems needed more than four partitions, a solution was found to go beyond the number of four. In the MBR, one partition could be created as an extended partition, as opposed to the other partitions that were created as primary partitions. Within the extended partition, multiple logical partitions could be created to reach a total number of 15 partitions that could be addressed by the Linux kernel.

Understanding the Need for GPT Partitioning

Current computer hard drives have become too big to be addressed by MBR partitions. That is one of the main reasons why a new partitioning scheme was needed. This partitioning scheme is the GUID Partition Table (GPT). On computers that are using the new Unified Extensible Firmware Interface (UEFI) as a replacement for the old BIOS system, GPT partitions are the only way to address disks. Also, older computer systems that are using BIOS instead of UEFI can be configured with GUID partitions, which is necessary if a disk with a size bigger than 2 TiB needs to be addressed.

Using GUID offers many benefits:

- The maximum partition size is 8 zebibyte (ZiB), which is $1024 \times 1024 \times 1024 \times 1024$ gibibytes.
- In GPT, up to a maximum number of 128 partitions can be created.
- The 2 TiB limit no longer exists.
- Because space that is available to store partitions is much bigger than 64 bytes, which was used in MBR, there is no longer a need to distinguish between primary, extended, and logical partitions.
- GPT uses a 128-bit global unique ID (GUID) to identify partitions.
- A backup copy of the GUID partition table is created by default at the end of the disk, which eliminates the single point of failure that exists on MBR partition tables.

Understanding Storage Measurement Units

When talking about storage, different measurement units are used. In some cases, units like megabyte (MB) are used. In other cases, units like mebibyte (MiB) are used. The difference between these two is that a megabyte is a multiple of 1,000, and a mebibyte is a multiple of 1,024. In computers, it makes sense to talk about multiples of 1,024 because that is how computers address items. However, confusion was created when hardware vendors a long time ago started referring to megabytes instead of mebibytes.

In the early days of computing, the difference was not that important. The difference between a kilobyte (KB) and a kibibyte (KiB) is just 24 bytes. The bigger the numbers grow, the bigger the difference becomes. A gigabyte, for instance, is $1,000 \times 1,000 \times 1,000$ bytes, so 1,000,000,000 bytes, whereas a gibibyte is $1,024 \times 1,024 \times 1,024$ bytes, which makes a total of 1,073,741,824 bytes, which is over 70 MB larger than 1 GB.

On current Linux distributions, the binary numbers (MiB, not MB) have become the standard—but do realize that some utilities still measure in MB and not MiB. In Table 14-2, you can see an overview of the values that are used.

In the past, KB, MB, and so on, were used both in decimal and binary situations; sometimes they were even mixed. For example, 1-Mbps line speed is one million bits per second. The once famous “1.44 MB” floppy disk was really 1,440,000 bytes in size (80 tracks $\times$ 2 heads $\times$ 9 sectors $\times$ 512-byte sectors), creating a mixed meaning of MB: $1.44 \times (\text{decimal K}) \times (\text{binary K})$.



Table 14-2 Disk Size Specifications

Symbol	Name	Value	Symbol	Name	Value
KB	Kilobyte	1000^1	KiB	Kibibyte	1024^1
MB	Megabyte	1000^2	MiB	Mebibyte	1024^2
GB	Gigabyte	1000^3	GiB	Gibibyte	1024^3
TB	Terabyte	1000^4	TiB	Tebibyte	1024^4
PB	Petabyte	1000^5	PiB	Pebibyte	1024^5
EB	Exabyte	1000^6	EiB	Exbibyte	1024^6
ZB	Zettabyte	1000^7	ZiB	Zebibyte	1024^7
YB	Yottabyte	1000^8	YiB	Yobibyte	1024^8

Managing Partitions and File Systems

As discussed in the previous section, there are two different types of partitions that can be used on RHEL 8. To match the different partition types, there are also two different partitioning utilities. The **fdisk** utility has been around for a long time and

is used to create MBR partitions. The **gdisk** utility is used to create GPT partitions. In this section, you learn how to use both.

Apart from **fdisk** and **gdisk**, there are other partitioning utilities as well, of which **parted** is probably the most important. Some people like it, as it is relatively easy to use, but at the same time it hides some of the more advanced features. For that reason, this chapter focuses on working with **fdisk** and **gdisk** and introduces **parted** only briefly.

For both MBR and GPT partitions, you need to specify the name of the disk device as an argument. Table 14-3 shows the most common disk device names that you work with on RHEL 8.

Table 14-3 Common Disk Device Types

Key
Topic

Device Name	Description
/dev/sda	A hard disk that uses the SCSI driver. Used for SCSI and SATA disk devices. Common on physical servers but also in VMware virtual machines.
/dev/nvme0n1	The first hard disk on an NVM Express (NVMe) interface. NVMe is a server-grade method to address advanced SSD devices. Note at the end of the device name that the first disk in this case is referred to as <i>n1</i> instead of <i>a</i> (as in common with the other types).
/dev/hda	The (legacy) IDE disk device type. You will seldom see this device type on modern computers.
/dev/sda	A disk in a KVM virtual machine that uses the virtio disk driver. This is the common disk device type for KVM virtual machines.
/dev/xvda	A disk in a Xen virtual machine that uses the Xen virtual disk driver. You see this when installing RHEL as a virtual machine in Xen virtualization. RHEL 8 cannot be used as a Xen hypervisor, but you might see RHEL 8 virtual machines on top of the Xen hypervisor using these disk types.

As you can see in Table 14-3, almost all disk device names end with the letter *a*. That is because it is the first disk that was found in your server. The second SCSI disk, for instance, would have the name `/dev/sdb`. If many disks are installed in a server, you can have up to `/dev/sdz` and even beyond. After `/dev/sdz`, the kernel continues creating devices with names like `/dev/sdaa` and `/dev/sdab`.

Creating MBR Partitions with fdisk

To create an MBR disk partition, you have to apply a multiple-step procedure, as shown in Exercise 14-1.

Exercise 14-1 Creating MBR Partitions with fdisk

This exercise has been written to use an installation of RHEL/CentOS that contains nonpartitioned disk space. If you do not have such an installation, you can use a second disk device on your demo environment. This can be a virtual disk that is added through your virtualization program, or a USB flash drive if you're working on a physical installation. In that case, make sure the device names in this exercise are replaced with the device names that match your hardware.

TIP In the end-of-chapter lab, you have to create partitions again. It will be a lot easier if at that point you can start from a clean installation. The following two steps help you revert easily to a system on which you have not created any partitions yet.

1. Type **dd if=/dev/sda of=/root/diskfile bs=1M count=1**. (If your disk is `/dev/vda` and not `/dev/sda`, change the disk name accordingly.) Using this command allows you to create a backup of the first megabyte of raw blocks and write that to the file `/root/diskfile`. This file allows you to easily revert to the situation that existed at the start of this exercise, as the command makes a backup of the MBR and all relevant metadata on disk.
2. Type **cp /etc/fstab /root/fstab** to make a backup of the `/etc/fstab` file as well.
3. Open a root shell and run the **fdisk** command. This command needs as its argument the name of the disk device where you want to create the partition. This exercise uses `/dev/sda`. Change that, if needed, according to your hardware.

```
[root@localhost ~]# fdisk /dev/sda
```

```
Welcome to fdisk (util-linux 2.32.1).
```

```
Changes will remain in memory only, until you decide to  
write them.
```

```
Be careful before using the write command.
```

```
Command (m for help):
```

4. Before doing anything, it is a good idea to check how much disk space you have available. Press **p** to see an overview of current disk allocation:

```
Command (m for help): p
```

```
Disk /dev/sda: 20 GiB, 21474836480 bytes, 41943040 sectors
```

```
Units: sectors of 1 * 512 = 512 bytes
```

```
Sector size (logical/physical): 512 bytes / 512 bytes
```

```
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: dos
Disk identifier: 0x7ad1a34b
```

```
Device      Boot      Start          End      Sectors   Size Id Type
/dev/sda1   *                2048    1026047    1024000   500M 83 Linux
/dev/sda2                1026048 24111103 23085056   11G 8e Linux LVM
```

In the output of this command, in particular look for the total number of sectors and the last sector that is currently used. If the last partition does not end on the last sector, you have available space to create a new partition.

5. Type `n` to add a new partition:

```
Command (m for help): n
```

```
Partition type
```

```
  p   primary (2 primary, 0 extended, 2 free)
  e   extended (container for logical partitions)
```

- 6.** Assuming you have a `/dev/sda1` and a `/dev/sda2` partition and nothing else, press **p** to create a primary partition. Accept the partition number that is now suggested, which should be `/dev/sda3`.
- 7.** Specify the first sector on disk that the new partition will start on. The first available sector is suggested by default, so press Enter to accept.
- 8.** Specify the last sector that the partition will end on. By default, the last sector available on disk is suggested. If you use that, after this exercise you will not have any disk space left to create additional partitions or logical volumes, so you should use another last sector. To use another last sector, you can do one of the following:
 - Enter the number of the last sector you want to use.
 - Enter **+number** to create a partition that sizes a specific number of sectors.
 - Enter **+number(K,M,G)** to specify the size you want to assign to the partition in KiB, MiB, or GiB.
 - Type **+1G** to make this a 1-GiB partition.

```
Command (m for help): n
```

```
Partition type
```

```
  p   primary (2 primary, 0 extended, 2 free)
  e   extended (container for logical partitions)
```

```
Select (default p): p
```

```
Partition number (3,4, default 3):
```

```
First sector (24111104-41943039, default 24111104):
```



```
Last sector, +sectors or +size{K,M,G,T,P} (24111104-41943039,
default 41943039): +1G
```

Created a new partition 3 of type 'Linux' and of size 1 GiB

After you enter the partition's ending boundary, **fdisk** will show a confirmation.

9. At this point, you can define the partition type. By default, a Linux partition type is used. If you want the partition to be of any other partition type, use **t** to change it. For this exercise there is no need to change the partition type. Common partition types include the following:

- **82:** Linux swap
- **83:** Linux
- **8e:** Linux LVM

10. If you are happy with the modifications, press **w** to write them to disk and exit **fdisk**. If you have created a partition on a disk that is already in use, you may now see the following message:

```
Command (m for help): w
```

```
The partition table has been altered!
```

```
Calling ioctl() to re-read partition table.
```

```
WARNING: Re-reading the partition table failed with error 16:
Device or resource busy.
```

```
The kernel still uses the old table. The new table will be
used at the next reboot or after you run partprobe(8) or
kpartx(8) Syncing disks.
```

```
[root@localhost ~]#
```

11. This message indicates that the partition has successfully been added to the partition table, but the in-memory kernel partition table could not be updated. You can see that by comparing the output of **fdisk -l /dev/sda** with the output of **cat /proc/partitions**, which shows the kernel partition table.
12. Type **partprobe /dev/sda** to write the changes to the kernel partition table. The partition has now been added, and you can create a file system on it as described in the section “Creating File Systems.”

NOTE You see the “re-reading the partition table failed with error 16” message only if you are adding partitions to a disk that already has some mounted partitions. If you are working on a new disk that does not have mounted partitions, you will not see this error, and you will not have to use the **partprobe** command.

Using Extended and Logical Partitions on MBR

In the previous procedure, you learned how to add a primary partition. If three partitions have been created already, there is room for one more primary partition, after which the partition table is completely filled up. If you want to go beyond four partitions on an MBR disk, you have to create an extended partition. Following that, you can create logical partitions within the extended partition.

Using logical partitions does allow you to go beyond the limitation of four partitions in the MBR; there is a disadvantage as well, though. All logical partitions exist within the extended partition. If something goes wrong with the extended partition, you have a problem with all logical partitions existing within it as well. If you need more than four separate storage allocation units, you might be better off using LVM instead of logical partitions. In Exercise 14-2 you learn how to work with extended and logical partitions.

NOTE An extended partition is only used for the purpose of creating logical partitions. You cannot create file systems directly on an extended partition!

Exercise 14-2 Creating Logical Partitions

1. In a root shell, type **fdisk /dev/sda** to open the **fdisk** interface.
2. To create a logical partition, when **fdisk** prompts which partition type you want to create, enter **e**.

```
Command (m for help): n
```

```
Partition type
```

```
p primary (3 primary, 0 extended, 1 free)
```

```
e extended (container for logical partitions)
```

```
Select (default e):
```

3. If the extended partition is the fourth partition that you are writing to the MBR, it will also be the last partition that can be added to the MBR. For that reason, it should fill the rest of your computer's hard disk. Press Enter to accept the default first sector and press Enter again when **fdisk** prompts for the last sector. Note that this extended partition should use the rest of the remaining disk space, because all disk space not included in this partition cannot be allocated at a later stage.

```
Using default response e.
```

```
Selected partition 4
```

```
First sector (26208256-41943039, default 26208256):
```

```
Last sector, +sectors or +size{K,M,G,T,P} (26208256-41943039,
default 41943039):
```

```
Created a new partition 4 of type 'Extended' and of size 7.5 GiB
```

4. Now that the extended partition has been created, you can create a logical partition within it. Still from the **fdisk** interface, press **n** again. The utility will prompt that all primary partitions are in use now and by default suggests adding a logical partition with partition number 5.

```
Command (m for help): n
All primary partitions are in use.
Adding logical partition 5
First sector (26210304-41943039, default 26210304):
```

5. Press Enter to accept the default first sector. When asked for the last sector, enter **+100M** (or any other size you want to use).

```
First sector (26210304-41943039, default 26210304):
Last sector, +sectors or +size{K,M,G,T,P} (26210304-41943039,
default 41943039): +100M
```

```
Created a new partition 5 of type 'Linux' and of size 100 MiB
```

6. Now that the logical partition has been created, enter **w** to write the changes to disk and quit **fdisk**. To complete the procedure, enter **partprobe** to update the kernel partition table. The new partition is now ready for use.

TIP The **fdisk** utility writes changes to disk only when you enter **w**, which is the **fdisk** write command. If you have made a mistake and want to get out, press **q** to quit.

In Exercise 14-2, you used the **partprobe** command to push changes in the partition table to the kernel partition table. This normally works out well, but in some cases does not. If at any time you are getting an error using **partprobe**, just reboot your computer, using the **reboot** command. If partitions have not been written to your system correctly, you really do not want to continue modifying and managing partitions, because you risk creating severe problems on your server.

Creating GPT Partitions with **gdisk**

If a disk is configured with a GUID Partition Table (GPT), or if it is a new disk that does not contain anything yet and has a size that goes beyond 2 TiB, you need to use the **gdisk** utility to create GUID partitions. This utility has a lot of similarities with **fdisk** but some differences as well. The following procedure shows how to create partitions in **gdisk**.

WARNING! Do not ever use **gdisk** on a disk that has been formatted with **fdisk** and already contains **fdisk** partitions. **gdisk** will detect that an MBR is present, and it will convert this to a GPT (see the following code listing). Your computer will most likely not be able to boot after doing this!

```
[root@localhost ~]# gdisk /dev/sda
GPT fdisk (gdisk) version 1.0.3
```

Partition table scan:

```
  MBR: MBR only
  BSD: not present
  APM: not present
  GPT: not present
```

```
*****
Found invalid GPT and valid MBR; converting MBR to GPT format in
memory. THIS OPERATION IS POTENTIALLY DESTRUCTIVE! Exit by typing 'q'
if you don't want to convert your MBR partitions to GPT format!
*****
```

Command (? for help):

To save you the hassle of going through this, I verified it does what it says. After converting an MBR to a GPT, your machine will not start anymore.

Exercise 14-3 demonstrates how to create partitions using **gdisk**.

Exercise 14-3 Creating GPT Partitions with **gdisk**

To apply the procedure in this exercise, you need a new disk device. Do *not* use a disk that contains data that you want to keep, because this exercise will delete all data on it. If you are using this exercise on a virtual machine, you may add the new disk through the virtualization software. If you are working on a physical machine, you can use a USB thumb drive as a disk device for this exercise. Note that this exercise works perfectly on a computer that starts from BIOS and not EFI; all you need is a dedicated disk device.

1. To create a partition with **gdisk**, type **gdisk /dev/sdb**. (Replace /dev/sdb with the exact device name used on your computer.) **gdisk** will try to detect the

current layout of the disk, and if it detects nothing, it will create the GPT partition table and associated disk layout.

```
[root@localhost ~]# gdisk /dev/sdb
GPT fdisk (gdisk) version 1.0.3
```

Partition table scan:

```
MBR: not present
BSD: not present
APM: not present
GPT: not present
```

Creating new GPT entries.

Command (? for help):

2. Type **n** to enter a new partition. You can choose any partition number between 1 and 128, but it is wise to accept the default partition number that is suggested.

```
Command (? for help): n
Partition number (1-128, default 1):
```

3. You now are asked to enter the first sector. By default, the first sector that is available on disk will be used, but you can specify an offset as well. This does not make sense, so just press Enter to accept the default first sector that is proposed.

```
First sector (34-2097118, default = 2048) or {+-}size{KMGTP}:
```

4. When asked for the last sector, by default the last sector that is available on disk is proposed (which would create a partition that fills the entire hard disk). You can specify a different last sector, or specify the disk size using **+**, the size, and KMGTP. So to create a 100-MiB disk partition, use **+100M**.

```
Last sector (2048-2097118, default = 2097118) or
{+-}size{KMGTP}: +100M
```

5. You now are asked to set the partition type. If you do not do anything, the partition type is set to 8300, which is the Linux file system partition type. Other options are available as well. You can press **l** to show a list of available partition types.

```
Current type is 'Linux filesystem'
Hex code or GUID (L to show codes, Enter = 8300): l
0700 Microsoft basic data  0c01 Microsoft reserved  2700 Windows RE
3000 ONIE boot              3001 ONIE config          3900 Plan 9
4100 PowerPC PReP boot      4200 Windows LDM data     4201 Windows LDM
                                metadata
4202 Windows Storage Spac  7501 IBM GPFS              7f00 ChromeOS
                                kernel
```

7f01 ChromeOS root	7f02 ChromeOS reserved	8200 Linux swap
8300 Linux filesystem	8301 Linux reserved	8302 Linux /home
8303 Linux x86 root (/)	8304 Linux x86-64 root (/	8305 Linux ARM64 root (/)
8306 Linux /srv	8307 Linux ARM32 root (/)	8400 Intel Rapid Start
8e00 Linux LVM	a000 Android bootloader	a001 Android bootloader 2
a002 Android boot	a003 Android recovery	a004 Android misc
a005 Android metadata	a006 Android system	a007 Android cache
a008 Android data	a009 Android persistent	a00a Android factory
a00b Android fastboot/ter	a00c Android OEM	a500 FreeBSD disklabel
a501 FreeBSD boot	a502 FreeBSD swap	a503 FreeBSD UFS
a504 FreeBSD ZFS	a505 FreeBSD Vinum/RAID	a580 Midnight BSD data
a581 Midnight BSD boot	a582 Midnight BSD swap	a583 Midnight BSD UFS
a584 Midnight BSD ZFS	a585 Midnight BSD Vinum	a600 OpenBSD disklabel
a800 Apple UFS	a901 NetBSD swap	a902 NetBSD FFS
a903 NetBSD LFS	a904 NetBSD concatenated	a905 NetBSD encrypted
a906 NetBSD RAID	ab00 Recovery HD	af00 Apple HFS/HFS+
af01 Apple RAID	af02 Apple RAID offline	af03 Apple label

The relevant partition types are as follows:

- **8200:** Linux swap
- **8300:** Linux file system
- **8e00:** Linux LVM

Notice that these are the same partition types as the ones that are used in MBR, with two 0s added to the IDs. You can also just press Enter to accept the default partition type 8300.

6. The partition is now created (but not yet written to disk). Press **p** to show an overview, which allows you to verify that this is really what you want to use.

Command (? for help): **p**

Disk /dev/sdb: 20971520 sectors, 10.0 GiB

Model: VMware Virtual S

```

Sector size (logical/physical): 512/512 bytes
Disk identifier (GUID): D9D1F660-E1DC-4737-8117-DC6990E665C9
Partition table holds up to 128 entries
Main partition table begins at sector 2 and ends at sector 33
First usable sector is 34, last usable sector is 20971486
Partitions will be aligned on 2048-sector boundaries
Total free space is 20766653 sectors (9.9 GiB)

Number   Start (sector)    End (sector)  Size      Code  Name
-----
1             2048           206847    100.0 MiB   8300   Linux
                                   filesystem

```

7. If you are satisfied with the current partitioning, press **w** to write changes to disk and commit. This gives a warning, after which the new partition table is written to the GUID partition table.

```
Command (? for help): w
```

```

Final checks complete. About to write GPT data. THIS WILL
OVERWRITE EXISTING
PARTITIONS!!

```

```
Do you want to proceed? (Y/N): Y
```

```
OK; writing new GUID partition table (GPT) to /dev/sdb.
```

```
The operation has completed successfully.
```

8. If at this point you get an error message indicating that the partition table is in use, type **partprobe** to update the kernel partition table.

Creating GPT Partitions with parted

As previously mentioned, apart from **fdisk** and **gdisk**, the **parted** utility can be used to create partitions. Because it lacks support for advanced features, I have focused on **fdisk** and **gdisk**, but I'd like to give you a quick overview of working with **parted**.

To use **parted**, you need to know that it has an interactive shell in which you can work with its different options. Exercise 14-4 guides you through the procedure of creating partitions using **parted**.

Exercise 14-4 Creating Partitions with parted

You need a new disk to work with this procedure. This exercise assumes that the new disk name is `/dev/sdc`.

1. From a root shell, type **parted /dev/sdc**. This opens the interactive **parted** shell.
2. Type **help** to get an overview of available commands.
3. Type **print**. You will see a message about an unrecognized disk label.
4. Type **mklabel** and press Enter. **parted** will now prompt for a disk label type. Press the Tab key twice to see a list of available disk label types. From the list, select **gpt** and press Enter.
5. Type **mkpart**. The utility prompts for a partition name. Type **part1** (the partition name doesn't really matter).
6. Now the utility prompts for a file system type. This is a very confusing option, because it suggests that you are setting a file system type here, but that is not the case. Also, when using Tab completion, you'll see a list of file systems that you've probably never used before. In fact, you could just press Enter to accept the default suggestion of `ext2`, as the setting isn't used anyway, but I suggest using a file system type that comes close to what you're going to use on the partition. So type **xfs** and press Enter to continue.
7. Now you are prompted for a start location. You can specify the start location as a number of blocks, or an offset from the start of the device. Notice that you can type **1M** to specify the start of the partition at 1 megabyte, or type **1 MiB** to have it start at 1 MiB. This is confusing, so make sure you specify the appropriate value here. At this point, type **1MiB** and press Enter.
8. Type **1GiB** to specify the end of the partition. After doing so, type **print** to print the current partition table, and type **quit** to quit the utility and commit your changes.

Creating File Systems

At this point, you know how to create partitions. A partition all by itself is not very useful. It only becomes useful if you decide to do something with it. That often means that you have to put a file system on top of it. In this section, you learn how to do that.

Different file systems can be used on RHEL 8. Table 14-4 provides an overview of the most common file systems.



Table 14-4 File System Overview

File System	Description
XFS	The default file system in RHEL 8.
Ext4	The default file system in previous versions of RHEL; still available and supported in RHEL 8.
Ext3	The previous version of Ext4. On RHEL 8, there is no need to use Ext3 anymore.
Ext2	A very basic file system that was developed in the early 1990s. There is no need to use this file system on RHEL 8 anymore.
BtrFS	A relatively new file system that is not supported in RHEL 8.
NTFS	A Windows-compatible file system that is not supported on RHEL 8.
VFAT	A file system that offers compatibility with Windows and Mac and is the functional equivalent of the FAT32 file system. Useful on USB thumb drives that exchange data with other computers but not on a server’s hard disks.

To format a partition with one of the supported file systems, you can use the **mkfs** command, using the option **-t** to specify which specific file system to use. Alternatively, one of the file system–specific tools can be used, such as **mkfs.ext4** to format an Ext4 file system.

NOTE If you use **mkfs** without any further specification of which file system you want to format, an Ext2 file system will be formatted. This is probably not what you want to use, so do not forget to specify which file system you want to use.

To format a partition with the default XFS file system, use the command **mkfs -t xfs**. Example 14-1 shows the output of this command.

Example 14-1 Formatting a File System with XFS

```
[root@localhost ~]# mkfs -t xfs /dev/sda5
meta-data=/dev/sda5          isize=512    agcount=4, agsize=6400 blks
                =             sectsz=512    attr=2, projid32bit=1
                =             crc=1        finobt=1, sparse=1, rmapbt=0
                =             reflink=1
data        =                bsize=4096    blocks=25600, imaxpct=25
                =                sunit=0     swidth=0 blks
naming      =version 2       bsize=4096    ascii-ci=0, ftype=1
log         =internal log    bsize=4096    blocks=1368, version=2
                =             sectsz=512    sunit=0 blks, lazy-count=1
realtime    =none           extsz=4096    blocks=0, rtextents=0
```

In Exercise 14-5, you learn how to create a file system.

Exercise 14-5 Creating a File System

In Exercise 14-1, you have created a partition `/dev/sda3`. In this exercise, you format it with an XFS file system. This exercise has one step only.

1. From a root shell, type **`mkfs.xfs /dev/sda3`**.

Changing File System Properties

When working with file systems, some properties can be managed as well.

File system properties are specific for the file system you are using, so you work with different properties and different tools for the different file systems.

Managing Ext4 File System Properties

The generic tool for managing Ext4 file system properties is **`tune2fs`**. This tool was developed a long time ago for the Ext2 file system and is compatible with Ext3 and Ext4 also. When managing Ext4 file system properties, **`tune2fs -l`** is a nice command to start with. Example 14-2 presents the output of this command where different file system properties are shown.

Example 14-2 Showing File System Properties with **`tune2fs -l`**

```
[root@localhost ~]# tune2fs -l /dev/sda3
tune2fs 1.44.3 (10-July-2018)
Filesystem volume name:   <none>
Last mounted on:         <not available>
Filesystem UUID:         ad8c3cf0-675b-4540-bc9a-4cfe63a90ffa
Filesystem magic number: 0xEF53
Filesystem revision #:    1 (dynamic)
Filesystem features:      has_journal ext_attr resize_inode dir_index
                           filetype extent 64bit flex_bg sparse_super
                           large_file huge_file dir_nlink extra_isize
                           metadata_csum
Filesystem flags:         signed_directory_hash
Default mount options:    user_xattr acl
Filesystem state:         clean
Errors behavior:          Continue
Filesystem OS type:       Linux
Inode count:              25688
Block count:              102400
```

```

Reserved block count:      5120
Free blocks:               93504
Free inodes:               25677
First block:               1
Block size:                1024
Fragment size:             1024
Group descriptor size:     64
Reserved GDT blocks:       256
Blocks per group:          8192
Fragments per group:       8192
Inodes per group:          1976
Inode blocks per group:    247
Flex block group size:     16
Filesystem created:        Wed Jun 19 03:29:06 2019
Last mount time:           n/a
Last write time:           Wed Jun 19 03:29:06 2019
Mount count:               0
Maximum mount count:       -1
Last checked:              Wed Jun 19 03:29:06 2019
Check interval:            0 (<none>)
Lifetime writes:           4441 kB
Reserved blocks uid:       0 (user root)
Reserved blocks gid:       0 (group root)
First inode:               11
Inode size:                128
Journal inode:             8
Default directory hash:    half_md4
Directory Hash Seed:       54b529bb-2ef1-4983-ae42-dfdef7ba8797
Journal backup:            inode blocks
Checksum type:             crc32c
Checksum:                  0xf36c81e5

```

As you can see, the **tune2fs -l** command shows many file system properties. One interesting property is the file system label, which shows as the Filesystem volume name. Labels are used to set a unique name for a file system, which allows the file system to be mounted in a consistent way, even if the underlying device name changes. Also interesting are the file system features and default mount options.

To change any of the default file system options, the **tune2fs** command enables you to do so with other parameters. Some common usage examples are listed here:

- Use **tune2fs -o** to set default file system mount options. When set to the file system, the option does not have to be specified while mounting through `/etc/fstab` anymore. Use, for instance, **tune2fs -o acl,user_xattr** to switch on access control lists and user-extended attributes. Use a `^` in front of the option to switch it off again, as in **tune2fs -o ^acl,user_xattr**.
- Ext file systems also come with file system features that may be enabled as a default. To switch on a file system feature, use **tune2fs -O** followed by the feature. To turn a feature off, use a `^` in front of the feature name.
- Use **tune2fs -L** to set a label on the file system. As described in the section “Mounting File Systems” later in this chapter, you can use a file system label to mount a file system based on its name instead of the device name. Instead of **tune2fs -L**, the **e2label** command enables you to do so.

Managing XFS File System Properties

The XFS file system is a completely different file system, and for that reason also has a completely different set of tools to manage its properties. It does not allow you to set file system attributes within the file system metadata. You can, however, change some XFS properties, using the **xfs_admin** command. For instance, use **xfs_admin -L mylabel** to set the file system label to mylabel.

Adding Swap Partitions

You use most of the partitions on a Linux server for regular file systems. On Linux, swap space is normally allocated on a disk device. That can be a partition or an LVM logical volume (discussed in Chapter 15). In case of an emergency, you can even use a file to extend the available swap space.

Using swap on Linux is a convenient way to improve Linux kernel memory usage. If a shortage of physical RAM occurs, non-recently used memory pages can be moved to swap, which makes more RAM available for programs that need access to memory pages. Most Linux servers for that reason are configured with a certain amount of swap. If swap starts being used intensively, you could be in trouble, though, and that is why swap usage should be closely monitored.

Sometimes, allocating more swap space makes sense. If a shortage of memory occurs, this shortage can be alleviated by allocating more swap space in some situations. (See Chapter 25, “Configuring Time Services,” for more information about system performance optimization.) This is done through a procedure where first a partition is created with the swap partition type, and then this partition is formatted as swap. Exercise 14-6 describes how to do this.

Exercise 14-6 Creating a Swap Partition

1. Type **fdisk /dev/sda** to open your disk in **fdisk**. (Use **gdisk** if you are using a disk with a GUID partition table.)
2. Press **n** to add a new partition. Specify start and stop cylinders and size.
3. Type **t** to change the partition type. If you are using **fdisk**, use partition type **82**. If you are using **gdisk**, use partition type **8200**.
4. Use **mkswap** to format the partition as swap space. Use, for instance, **mkswap /dev/sda6** if the partition you have just created is **/dev/sda6**.
5. Type **free -m**. You see the amount of swap space that is currently allocated. This does not include the swap space you have just created, as it still needs to be activated.
6. Use **swapon** to switch on the newly allocated swap space. If, for instance, the swap device you have just created is **/dev/sda6**, use **swapon /dev/sda6** to activate the swap space.
7. Type **free -m** again. You see that the new swap space has been added to your server.

Adding Swap Files

If you do not have free disk space to create a swap partition and you do need to add swap space urgently, you can use a swap file as well. From a performance perspective, it does not even make that much difference if a swap file is used instead of a swap device such as a partition or a logical volume, and it may help you fulfill an urgent need in a timely manner.

To add a swap file, you need to create the file first. The **dd if=/dev/zero of=/swapfile bs=1M count=100** command would add 100 blocks with a size of 1 MiB from the **/dev/zero** device (which generates 0s) to the **/swapfile** file. The result is a 100-MiB file that can be configured as swap. To do so, you can follow the same procedure as for swap partitions. First use **mkswap /swapfile** to mark the file as a swap file, and then use **swapon /swapfile** to activate it.

Mounting File Systems

Just creating a partition and putting a file system on it is not enough to start using it. To use a partition, you have to mount it as well. By mounting a partition (or better, the file system on it), you make its contents accessible through a specific directory.

To mount a file system, some information is needed:

- **What to mount:** This information is mandatory and specifies the name of the device that needs to be mounted.
- **Where to mount it:** This is also mandatory information that specifies the directory on which the device should be mounted.
- **What file system to mount:** Optionally, you can specify the file system type. In most cases, this is not necessary. The **mount** command will detect which file system is used on the device and make sure the correct driver is used.
- **Mount options:** Many mount options can be used when mounting a device. Using options is optional and depends on the needs you may have for the file system.

Manually Mounting File Systems

To manually mount a file system, the **mount** command is used. To disconnect a mounted file system, the **umount** command is used. Using these commands is relatively easy. To mount the file system that is on `/dev/sda5` on the directory `/mnt`, for example, use the following command:

```
mount /dev/sda5 /mnt
```

To disconnect the mount, you can use **umount** with either the name of the device or the name of the mount point you want to disconnect. So, both of the following commands will work:

```
umount /dev/sda5  
umount /mnt
```

Using Device Names, UUIDs, or Disk Labels

To mount a device, the name of the device can be used, as in the command **mount /dev/sda5 /mnt**. If your server is used in an environment where a dynamic storage topology is used, this is not always the best approach. You may today have a storage device `/dev/sda5`, which after changes in the storage topology can be `/dev/sdb5` after the next reboot of your server. This is why on a default RHEL 8 installation, UUIDs are used instead of device names.

Every file system by default has a UUID associated to it—not just file systems that are used to store files but also special file systems such as the swap file system. You can use the **blkid** command to get an overview of the current file systems on your system and the UUID that is used by that file system.

Before the use of UUIDs was common, file systems were often configured to work with labels, which can be set using the **e2label** command or the **xfs_admin -L** command. This has become more uncommon in recent Linux versions. If a file system has a label, the **blkid** command will also show it, as can be seen in Example 14-3.

Example 14-3 Using **blkid** to Find Current File System UUIDs

```
[root@server3 ~]# blkid
/dev/sda1: UUID="02305166-840d-4f74-a868-c549b3229e65" TYPE="xfs"
/dev/sda2: UUID="Hasz5q-nZF3-XN94-L2fm-xUwz-3VEq-qVCAYi"
          TYPE="LVM2_member"
/dev/sda5: UUID="42f419c4-633f-4ed7-b161-519a4dadd3da" TYPE="xfs"
/dev/mapper/centos-swap: UUID="5867ba02-fd89-475c-be56-7922febde43b"
                       TYPE="swap"
/dev/mapper/centos-root: UUID="b2022ac4-73c6-4e6b-a52f-703e3e2476b7"
                        TYPE="xfs"
```

To mount a file system based on a UUID, you use **UUID=nnnnnn** instead of the device name. So if you want to mount `/dev/sda5` from Example 14-3 based on its UUID, the command becomes as follows:

```
mount UUID="42f419c4-633f-4ed7-b161-519a4dadd3da" /mnt
```

Manually mounting devices using the UUID is not exactly easier. If mounts are automated as discussed in the next section, however, it does make sense using UUIDs instead of device names.

To mount a file system using a label, you use the **mount LABEL=labelname** command. For example, use **mount LABEL=mylabel /mnt** to temporarily mount the file system with the name `mylabel` on the `/mnt` directory.

Automating File System Mounts Through `/etc/fstab`

Normally, you do not want to be mounting file systems manually. Once you are happy with them, it is a good idea to have them mounted automatically. The classical way to do this is through the `/etc/fstab` file. Example 14-4 shows sample contents of this file.

Example 14-4 Sample /etc/fstab File Contents

```
[root@server3 ~]# cat /etc/fstab

#
# /etc/fstab
# Created by anaconda on Fri Jan 16 10:28:41 2015
#
# Accessible filesystems, by reference, are maintained under
# '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8)
# for more info
#
/dev/mapper/centos-root /          xfs     defaults 1 1
UUID=02305166-840d-4f74 /          xfs     defaults 1 2
/dev/mapper/centos-swap swap      swap    defaults 0 0
```

In the /etc/fstab file, everything is specified to mount the file system automatically. For this purpose, every line has six fields, as summarized in Table 14-5.

Table 14-5 /etc/fstab Fields

Key
Topic

Field	Description
Device	The device that must be mounted. A device name, UUID, or label can be used.
Mount Point	The directory or kernel interface where the device needs to be mounted.
File System	The file system type.
Mount Options	Mount options.
Dump Support	Use 1 to enable support to back up using the dump utility. This may be necessary for some backup solutions.
Automatic Check	Specifies whether the file system should be checked automatically when booting. Use 0 to disable automated check, 1 if this is the root file system and it has to be checked automatically, and 2 for all other file systems that need automatic checking while booting. Network file systems should have this option set to 0.

Based on what has previously been discussed about the **mount** command, you should have no problem understanding the Device, Mount Point, and File System fields in /etc/fstab. Notice that in the mount point not all file systems use a directory name. Some system devices such as swap are not mounted on a directory, but on a kernel interface. It is easy to recognize when a kernel interface is used; its name does not start with a / (and does not exist in the file system on your server).

The Mount Options field defines specific mount options that can be used. If no specific options are required, this line will just read “defaults.” To offer specific functionality, a large number of mount options can be specified here. Table 14-6 gives an overview of some of the more common mount options.

**Key
Topic**
Table 14-6 Common Mount Options

Option	Use
auto/ noauto	The file system will [not] be mounted automatically.
acl	Adds support for file system access control lists (see Chapter 7, “Permissions Management”).
user_xattr	Adds support for user-extended attributes (see Chapter 7).
ro	Mounts the file system in read-only mode.
atime / noatime	Disables or enables access time modifications.
noexec / exec	Denies or allows execution of program files from the file system.
_netdev	Use this to mount a network file system. This tells fstab to wait until the network is available before mounting this file system.

The fifth column of `/etc/fstab` specifies support for the dump utility. This is a utility that was developed to create file system backups. It is good practice to switch this feature on by specifying a 1 for all real file systems, and switch it off by specifying 0 for all system mounts.

The last column indicates if the file system integrity needs to be checked while booting. Put a 0 if you do not want to check the file system at all, a 1 if this is the root file system that needs to be checked before anything else, and a 2 if this is a nonroot file system that needs to be checked while booting.

In Exercise 14-7, you learn how to mount partitions through `/etc/fstab` by mounting the XFS-formatted partition `/dev/sda5` that you created in previous exercises.

WARNING If a file system through `/etc/fstab` is flagged for automatic file system check and something prevents the file system from being checked correctly, your system stops booting and prompts “enter root password to enter maintenance mode.” To prevent this from ever happening, you could choose to disable automated checks while booting. See Chapter 18, “Essential Troubleshooting Skills,” for more information on how to fix this specific case.

Exercise 14-7 Mounting Partitions Through /etc/fstab

1. From a root shell, type **blkid**. Use the mouse to copy the `UUID="nnnn"` part for `/dev/sda5`.
 2. Type **mkdir -p /mounts/data** to create a mount point for this partition.
 3. Open `/etc/fstab` in an editor and add the following line:
`UUID="nnnn" /mounts/data xfs defaults 1 2`
 4. Before attempting an automatic mount while rebooting, it is a good idea to test the configuration. Type **mount -a**. This mounts everything that is specified in `/etc/fstab` and that has not been mounted already.
 5. Type **df -h** to verify that the partition has been mounted correctly.

Summary

In this important chapter, you learned how to work with partitions and file systems on RHEL 8. You learned how to create partitions for MBR and GPT disks and how to put a file system on top of the partition. You also learned how to mount these partitions manually and automatically through `/etc/fstab`.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 14-7 lists a reference of these key topics and the page numbers on which each is found.

Table 14-7 Key Topics for Chapter 14

Key Topic Element	Description	Page
Table 14-2	Disk size specifications	314
Table 14-3	Common disk device types	315
Table 14-4	File system types	326
Table 14-5	<code>/etc/fstab</code> fields	333
Table 14-6	Common mount options	334



Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

BIOS, MBR, partition, primary partition, extended partition, logical partition, GPT, mount, umount, UUID, label, Ext2, Ext3, Ext4, XFS, BtrFS, VFAT, fstab

Review Questions

The questions that follow use an open-ended format that is meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. Which tool do you use to create GUID partitions?
2. Which tool do you use to create MBR partitions?
3. What is the default file system on RHEL 8?
4. What is the name of the file that is used to automatically mount partitions while booting?
5. Which mount option do you use if you want a file system not to be mounted automatically while booting?
6. Which command enables you to format a partition that has type 82 with the appropriate file system?
7. You have just added a couple of partitions for automatic mounting while booting. How can you safely test if this is going to work without actually rebooting?

8. Which file system is created if you use the **mkfs** command without any file system specification?
9. How do you format an Ext4 partition?
10. How do you find UUIDs for all devices on your computer?

End-of-Chapter Lab

In the exercises you have worked through in this chapter, you have already created partitions. Before working on this end-of-chapter lab, it is a good idea to start from a clean installation. In Exercise 14-1, you created some backup files. Before starting to work on this lab, restore the original setup using the following two steps:

1. Type **dd if=/dev/diskfile of=/dev/sda**. (Use **of=/dev/vda** if your disk device is **/dev/vda** instead of **/dev/sda**.)
2. Copy the backup of the **/etc/fstab** file, using **cp /root/fstab /etc**.

This restores the original configuration. You are now ready to start the end-of-chapter lab. After successfully completing this lab, repeat this procedure. This allows you to work on clean disks when creating LVM logical volumes, as described in the next chapter.

Lab 14.1

1. Add two partitions to your server. If possible, put them on the primary disk that is in use on your server. If that is not possible, use a second (virtual or USB) disk to add these partitions. Create both partitions with a size of 100 MiB. One of these partitions must be configured as swap space; the other partition must be formatted with an Ext4 file system.
2. Configure your server to automatically mount these partitions. Mount the Ext4 partition on **/mounts/data** and mount the swap partition as swap space.
3. Reboot your server and verify that all is mounted correctly. In case of problems, read Chapter 18, “Essential Troubleshooting Skills,” for tips on how to troubleshoot.



The following topics are covered in this chapter:

- Understanding LVM
- Creating LVM Logical Volumes
- Resizing LVM Logical Volumes
- Configuring Stratis
- Configuring VDO

The following RHCSA exam objectives are covered in this chapter:

- Create and remove physical volumes
- Assign physical volumes to volume groups
- Create and delete logical volumes
- Extend existing logical volumes
- Configure disk compression
- Manage layered storage

In Chapter 14, “Managing Storage,” you learned how to manage partitions on a hard disk. Creating multiple partitions on a disk is useful because it enables you to keep different data types in separate partitions, but it does not offer the flexibility that the advanced storage solutions are offering. In this chapter, you’ll learn how to work with advanced storage solutions, including Logical Volume Manager (LVM), Stratis, and Virtual Data Optimizer (VDO).

Managing Advanced Storage

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 15-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 15-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Understanding LVM	1–2
Creating LVM Logical Volumes	3–5
Resizing LVM Logical Volumes	6
Configuring Stratis	7–8
Configuring VDO	9–10

1. Which of the following is not a standard component in an LVM setup?
 - a. Logical volume
 - b. File system
 - c. Volume group
 - d. Physical volume
2. Which of the following is not an LVM feature?
 - a. Volume resizing
 - b. Hot replacement of failing disk
 - c. Copy on write
 - d. Snapshots

3. Which partition type do you need on a GPT partition to mark it with the LVM partition type?
 - a. 83
 - b. 8e
 - c. 8300
 - d. 8e00
4. Which of the following commands shows correctly how to create a logical volume that uses 50% of available disk space in the volume group?
 - a. **vgadd -n lvdata -l +50%FREE vgdata**
 - b. **lvcreate lvdata -l 50%FREE vgdata**
 - c. **lvcreate -n lvdata -l 50%FREE vgdata**
 - d. **lvadd -n lvdata -l 50% FREE /dev/vgdata**
5. Which command shows an overview of available physical volumes? (Choose two.)
 - a. **pvshow**
 - b. **pvdisplay**
 - c. **pvs**
 - d. **pvlist**
6. Which statement about resizing LVM logical volumes is *not* true?
 - a. The Ext4 file system can be increased and decreased in size.
 - b. Use **lvextend** with the **-r** option to automatically resize the file system.
 - c. The XFS file system cannot be resized.
 - d. To increase the size of a logical volume, you need allocatable space in the volume group.
7. How much storage is used in a Stratis file system for metadata storage?
 - a. 527 MiB
 - b. 1 GiB
 - c. 4 MiB
 - d. 4 GiB

8. Which of the following lines correctly shows how a Stratis file system should be mounted through `/etc/fstab`?
- a. `UUID=abcd /stratis xfs defaults 0 0`
 - b. `/dev/stratis/stratis1 /stratis xfs defaults,x-systemd.requires=stratis.service 0 0`
 - c. `UUID=abcd /stratis xfs defaults,x-systemd.requires=stratis.service 0 0`
 - d. `/dev/stratis/stratis1 /stratis xfs defaults 0 0`
9. Which of the following environments could benefit from using VDO?
(Choose all that apply.)
- a. Virtualization
 - b. Databases
 - c. Containers
 - d. Cloud storage
10. Which option should you use while formatting a VDO device with the XFS file system to ensure that empty blocks are not immediately discarded?
- a. `-k`
 - b. `-K`
 - c. `-d`
 - d. `-D`

Foundation Topics

Understanding LVM

In the early days of Linux servers, storage was handled by creating partitions on disks. Even if this approach does work, there are some disadvantages, the most important of which is that disks are inflexible. That is why the Logical Volume Manager was introduced. Whereas it is not possible to dynamically grow a partition that is running out of disk space, this is possible when working with LVM. LVM offers many other advantages as well, which you learn about in this chapter.

LVM Architecture

In the LVM architecture, several layers can be distinguished. On the lowest layer, the storage devices are used. These can be any storage devices, such as complete disks, partitions, logical units (LUNs) on a storage-area network (SAN), and whatever else is made possible in modern storage topologies. In this chapter you'll learn how to use partitions as physical volumes, which is recommended practice. By using partitions instead of complete disk devices, it is easy for other tools to recognize that some storage has already been configured on the block device, which makes it less likely that misconfigurations are going to occur.


The storage devices need to be flagged as physical volumes, which makes them usable in an LVM environment and makes them usable by other utilities trying to gain access to the logical volume. A storage device that is a physical volume can be added to the volume group, which is the abstraction of all available storage. The “abstraction” means that the volume group is not something that is fixed, but it can be resized when needed, which makes it possible to add more space on the volume group level when volumes are running out of disk space. The idea is simple: If you are running out of disk space on a logical volume, you take available disk space from the volume group. And if there is no available disk space in the volume group, you just add it by adding a physical volume.

On top of the volume group are the logical volumes. Logical volumes do not act on disks directly but get their disk space from available disk space in the volume group. That means that a logical volume may consist of available storage from multiple physical volumes, which adds an important layer of additional flexibility to the storage configuration.

NOTE It is a good idea to avoid logical volumes from spanning multiple physical volumes; if one of the physical volumes breaks, all files on the LVM file system will become inaccessible.

The actual file systems are created on the logical volumes. Because the logical volumes are flexible with regard to size, that makes the file systems flexible as well. If a file system is running out of disk space, it is relatively easy to extend the file system or to reduce it if the file system allows that. Note that in order to resize file systems when logical volumes are resized, the file systems must offer support for that.

Figure 15-1 gives an overview of the LVM architecture.



**Key
Topic**

FIGURE 15-1 LVM Architecture Overview

LVM Features

There are several reasons why LVM is great. The most important reason is that LVM offers a flexible solution for managing storage. Volumes are no longer bound to the restrictions of physical hard drives. If additional storage space is needed, the volume group can easily be extended by adding a new physical volume, so that disk space can be added to the logical volumes. It is also possible to reduce the size of a logical volume, but only if the file system that was created on that volume supports the feature of reducing the size of the file system.

Another important reason why administrators like using LVM is the support for snapshots. A *snapshot* keeps the current state of a logical volume and can be used to revert to a previous situation or to make a backup of the file system on the logical volume if the volume is open. Using snapshots is essential in backup strategies.

LVM snapshots are created by copying the logical volume administrative data (the metadata) that describes the current state of files to a snapshot volume. As long as nothing changes, from the LVM snapshot metadata the original blocks in the original volume are addressed. When blocks are modified, the blocks containing the previous state of the file are copied over to the snapshot volume, which for that reason will grow. Using this method ensures that, by accessing an LVM snapshot volume, the exact state of the files as they were when the snapshot was created can be accessed. Because the snapshot will grow when files on the original volume change, when planning for snapshots, you should make sure that a sufficient amount of disk space is available. Also, snapshots are supposed to be temporary: once a snapshot has served its purpose, it can be removed.

A third important advantage of using LVM logical volumes is the option to replace failing hardware easily. If a hard disk is failing, data can be moved within the volume group (through the **pvmove** command), the failing disk can then be removed from the volume group, and a new hard disk can be added dynamically, without requiring any downtime for the logical volume itself.

Creating LVM Logical Volumes

Creating LVM logical volumes involves creating the three layers in the LVM architecture. You first have to convert physical devices, such as disks or partitions, into physical volumes (PVs); then you need to create the volume group (VG) and assign PVs to it. As the last step, you need to create the logical volume (LV) itself. In this section, you learn what is involved in creating these three layers.

Different utilities exist for creating LVM logical volumes. This chapter focuses on using the command-line utilities. They are relatively easy to use, and they are available in all environments (whether you are running a graphical interface or not).

TIP You absolutely do not need to memorize the commands discussed in this chapter for the RHCSA exam. All you really need to remember is **pv**, **vg**, and **lv**. Open a command line, type **pv**, and press the Tab key twice. This will show all commands that start with *pv*, which are all commands that are used for managing physical volumes. After you have found the command you need, run this command with the **--help** option. This shows a usage summary that lists everything you must do to create the element you need. Example 15-1 shows an example of the **pvcreate --help** command (which is explained in the next subsection).

Example 15-1 Requesting Help for the **pvcreate** Command

```
[root@control ~]# pvcreate --help
pvcreate - Initialize physical volume(s) for use by LVM

pvcreate PV ...
    [ -f|--force ]
    [ -M|--metadatatype lvm2 ]
    [ -u|--uuid String ]
    [ -Z|--zero y|n ]
    [     --dataalignment Size[k|UNIT] ]
    [     --dataalignmentoffset Size[k|UNIT] ]
    [     --bootloaderareasize Size[m|UNIT] ]
    [     --labelsector Number ]
    [     --pvmetadatacopies 0|1|2 ]
    [     --metadatasize Size[m|UNIT] ]
    [     --metadataignore y|n ]
    [     --norestorefile ]
    [     --setphysicalvolumesize Size[m|UNIT] ]
    [     --reportformat basic|json ]
    [     --restorefile String ]
    [ COMMON_OPTIONS ]

Common options for lvm:
    [ -d|--debug ]
    [ -h|--help ]
    [ -q|--quiet ]
    [ -v|--verbose ]
    [ -y|--yes ]
    [ -t|--test ]
    [     --commandprofile String ]
    [     --config String ]
    [     --driverloaded y|n ]
    [     --nolocking ]
    [     --lockopt String ]
    [     --longhelp ]
    [     --profile String ]
    [     --version ]

Use --longhelp to show all options and advanced commands.
```

Creating the Physical Volumes

Before you can use the LVM tools to create physical volumes, you need to create a partition marked as the LVM partition type. This is basically the same procedure as described in Chapter 14, with the only difference being that before writing changes to disk, you need to change the partition type.

In **fdisk** and **gdisk**, you can use **t** from the menu to change the type. If you are using an MBR disk, the partition type is 8e. If you are using a GUID disk, use the partition type 8e00. If you are using **parted** to create partitions, you need to use the **set *n* lvm on** command from within the **parted** interface (where *n* is the number of the partition you want to mark for use with LVM).

After creating the partition and flagging it as an LVM partition type, you need to use **pvccreate** to mark it as a physical volume. This writes some metadata to the partition, which allows it to be used in a volume group. The entire procedure is summarized in Exercise 15-1, in which you create a physical volume. Also see Example 15-2 for an overview of this procedure.

Exercise 15-1 Creating a Physical Volume

To do this exercise, you need a hard disk that has free (unpartitioned) disk space available. The recommended method to make disk space available is by adding a new hard disk in your virtual machine environment. In this exercise, I use a clean `/dev/sdc` device to create the partition. You may have to change the device name to match your configuration. If you do not have a dedicated hard disk available to create this configuration, you might want to consider attaching a USB key to your machine.

1. Open a root shell and type **parted /dev/sdc**.
2. Type **print**. This will show the current partition table layout. There should be none at this point.
3. Type **mklabel msdos** to set the MBR-compatible partition type.
4. Type **mkpart** to start the procedure to create a partition, and enter **primary** when asked for the partition type.
5. Type **xfs** to specify that the XFS file system should be used. When asked for the Start position of the partition, type **1MiB**, and when asked for the end, enter **1GiB**.
6. Type **set 1 lvm on** to enable the LVM partition type on the partition.
7. Now that the partition with the correct partition type has been created, type **quit** to close the **parted** interface.

Example 15-2 Creating an LVM Partition in **parted**

```
[root@server1 ~]# parted /dev/sdc
GNU Parted 3.2
Using /dev/sdc
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) print
Error: /dev/sdc: unrecognised disk label
Model: VMware, VMware Virtual S (scsi)
Disk /dev/sdc: 21.5GB
Sector size (logical/physical): 512B/512B
Partition Table: unknown
Disk Flags:
(parted) mklabel msdos
(parted) mkpart
Partition type? primary/extended? primary
File system type? [ext2]? xfs
Start? 1MiB
End? 1GiB
(parted) set 1 lvm on
(parted) print
Model: VMware, VMware Virtual S (scsi)
Disk /dev/sdb: 21.5GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

Number  Start   End     Size    Type    File system  Flags
  1      1049kB  1074MB  1073MB  primary xfs          lvm, lba

(parted) quit
Information: You may need to update /etc/fstab.
```

8. Now that the partition has been created, you need to flag it as an LVM physical volume. To do this, type **pvcreeate /dev/sdc1**. You should now get this output: Physical volume “/dev/sdc1” successfully created.
9. Type **pvs** to verify that the physical volume has been created successfully. The output may look like Example 15-3. Notice that in this listing another physical volume already exists; that is because RHEL uses LVM by default to organize storage.

Example 15-3 Verifying the Physical Volume

```
[root@server1 ~]# pvs
PV          VG      Fmt  Attr PSize    PFree
/dev/sda2   rhel   lvm2 a--   <19.00g    0
/dev/sdc1           lvm2 ---  1023.00m 1023.00m
```

As an alternative to the **pvs** command, which shows a summary of the physical volumes and their attributes, you can also use the **pvdisplay** command to show more details. Example 15-4 shows an example of the output of this command.

Example 15-4 Example **pvdisplay** Command Output

```
[root@server1 ~]# pvdisplay
--- Physical volume ---
PV Name                /dev/sda2
VG Name                 rhel
PV Size                 <19.00 GiB / not usable 3.00 MiB
Allocatable             yes (but full)
PE Size                 4.00 MiB
Total PE                 4863
Free PE                  0
Allocated PE             4863
PV UUID                 h9seBL-Z0AK-05xH-RNjO-Ui4d-d8Px-VunuQJ

"/dev/sdb1" is a new physical volume of "1023.00 MiB"
--- NEW Physical volume ---
PV Name                /dev/sdc1
VG Name
PV Size                 1023.00 MiB
Allocatable             NO
PE Size                 0
Total PE                 0
Free PE                  0
Allocated PE             0
PV UUID                 vMqYr7-70nM-kNeW-am8J-c9aA-VWUF-vxw0XJ
```

If you want a very compact overview of the current storage configuration on your server, you might also like the **lsblk** command. As shown in Example 15-5, this command gives a hierarchical overview of which disks and partitions are used in what LVM volume groups and logical volumes.

Example 15-5 Using **lsblk** for a Compact Overview of the Current Storage Configuration

```
[root@server1 ~]# lsblk
```

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
sda	8:0	0	20G	0	disk	
└─sda1	8:1	0	1G	0	part	/boot
└─sda2	8:2	0	19G	0	part	
└─rhel-root	253:0	0	17G	0	lvm	/
└─rhel-swap	253:1	0	2G	0	lvm	[SWAP]
sdc	8:16	0	20G	0	disk	
└─sdc1	8:17	0	1023M	0	part	
sdb	8:32	0	20G	0	disk	
sr0	11:0	1	6.6G	0	rom	/run/media/student/ RHEL-8-0-0-BaseOS-x86_64

Creating the Volume Groups

Now that the physical volume has been created, you can assign it to a volume group. It is possible to add a physical volume to an existing volume group (which is discussed later in this chapter), but you will now learn how to create a new volume group and add the physical volume to it. This is a simple one-command procedure. Just type **vgcreate** followed by the name of the volume group you want to create and the name of the physical device you want to add to it. So, if the physical volume name is `/dev/sdc1`, the complete command is **vgcreate vgdata /dev/sdc1**. You are completely free in your choice of name for the volume group. I like to start all volume group names with *vg*, which makes it easy to find the volume groups if there are many, but you are free to choose anything you like.

Between the previous section and the preceding paragraph, you have learned how to create a volume group in a two-step procedure where you first create the physical volume with the **pvcreate** command and then add the volume group using the **vgcreate** command. You can do this in a one-step procedure as well (where using a separate **pvcreate** command will not be necessary). If you are adding a partition to the volume group, however, it must be marked as an LVM partition type already.

The one-step procedure is particularly useful for adding a complete disk device (which does not need to be marked as anything). If you want to add the disk `/dev/sdc`, for instance, just type **vgcreate vgdata /dev/sdc** to create a volume group `vgdata` that contains the `/dev/sdc` device. When you are doing this to add a device that has not been marked as a physical volume yet, the **vgcreate** utility will automatically flag it as a physical volume.

When creating volume groups, a physical extent size is used. The physical extent size defines the size of the building blocks used to create logical volumes. A logical volume always has a size that is a multiple of the physical extent size. If you need to create huge logical volumes, it is more efficient to use a big physical extent size. If you do not specify anything, a default extent size of 4.00 MiB is used. The physical extent size is always specified as a multiple of 2 MiB, with a maximum size of 128 MiB. Use the **vgcreate -s** option to specify the physical extent size you want to use.

NOTE When working with LVM, there is the *physical* extent size to consider. This is the size of the basic building blocks used in the LVM configuration. When working with an ext4 file system, *logical* extents are used. The extent sizes on LVM are in no way related to the extent sizes that are used on the file systems.

After creating the volume group, you can request details about the volume group using the **vgs** command for a short summary, or the **vgdisplay** command to get more information. Example 15-6 shows an example of the output of the **vgdisplay** command.

Example 15-6 Showing Current Volume Group Properties

```
[root@server1 ~]# vgdisplay
...
--- Volume group ---
VG Name                vgdata
System ID
Format                 lvm2
Metadata Areas         1
Metadata Sequence No   2
VG Access               read/write
VG Status               resizable
MAX LV                 0
Cur LV                 1
Open LV                 0
Max PV                 0
Cur PV                 1
Act PV                 1
VG Size                 1020.00 MiB
PE Size                 4.00 MiB
Total PE                255
Alloc PE / Size        127 / 508.00 MiB
Free PE / Size          128 / 512.00 MiB
VG UUID                b68TU5-UPQv-xQRB-7Djx-SuZf-agDb-vAuJYD
```

Creating the Logical Volumes and File Systems

Now that the volume group has been created, you can start creating logical volumes from it. This procedure is slightly more complicated than the creation of physical volumes or volume groups because there are more choices to be made. While creating the logical volume, you must specify a volume name and a size.

The volume size can be specified as an absolute value using the **-L** option. Use, for instance, **-L 5G** to create an LVM volume with a 5-GiB size. Alternatively, you can use relative sizes with the **-l** option. For instance, use **-l 50%FREE** to use half of all available disk space. You'll further need to specify the name of the volume group that the logical volume is assigned to, and optionally (but highly recommended), you can use **-n** to specify the name of the logical volume. For instance, use **lvcreate -n lvvol1 -L 100M vgdata** to create a logical volume with the name `lvvol1` and add that to the `vgdata` volume group. Once the logical volume has been created, you can use the **mkfs** utility to create a file system on top of it.

Understanding LVM Device Naming

Now that the logical volume has been created, you can start using it. To do this, you need to know the device name. LVM volume device names can be addressed in multiple ways. The simple method is to address the device as `/dev/vgname/lvname`. So, for example, if you have created a volume with the name `lvdata`, which gets its available disk space from the `vgdata` volume group, the device name would be `/dev/vgdata/lvdata`, which is in fact a symbolic link to the device mapper name (which is explained next).

For naming LVM volumes, another system plays a role: device mapper. The device mapper (abbreviated as `dm`) is a generic interface that the Linux kernel uses to address storage devices. Device mapper is used by multiple device types, such as LVM volumes, but also by software RAID and advanced network devices such as multipath devices.

Device mapper devices are generated on detection and use meaningless names like `/dev/dm-0` and `/dev/dm-1`. To make these devices easier to access, device mapper creates symbolic links in the `/dev/mapper` directory that point to these meaningless device names. The symbolic links follow the naming structure `/dev/mapper/vgname-lvname`.

So, the device `/dev/vgdata/lvdata` would also be known as `/dev/mapper/vgdata-lvdata`. When working with LVM logical volumes, you can use either of these device names. Example 15-7 shows an overview of the different LVM device names as provided by the device mapper. In Exercise 15-2, you learn how to create a volume group and logical volumes.

Example 15-7 LVM Device Name Overview

```
[root@server1 ~]# \ls -l /dev/vgdata/lvdata
lrwxrwxrwx. 1 root root 7 May 19 06:46 /dev/vgdata/lvdata -> ../dm-2
[root@server1 ~]# \ls -l /dev/mapper/
total 0
crw-----. 1 root root 10, 236 May 19 06:33 control
lrwxrwxrwx. 1 root root      7 May 19 06:33 rhel-root -> ../dm-0
lrwxrwxrwx. 1 root root      7 May 19 06:33 rhel-swap -> ../dm-1
lrwxrwxrwx. 1 root root      7 May 19 06:46 vgdata-lvdata -> ../dm-2
```

Exercise 15-2 Creating the Volume Group and Logical Volumes

In Exercise 15-1, you created a physical volume. In this exercise, you continue working on that physical volume and assign it to a volume group. Then you add a logical volume from that volume group. You can work on this exercise only after successful completion of Exercise 15-1.

1. Open a root shell. Type **pvs** to verify the availability of physical volumes on your machine. You should see the **/dev/sdc1** physical volume that was created previously.
2. Type **vgcreate vgdata /dev/sdc1**. This will create the volume group with the physical volume assigned to it.
3. Type **vgs** to verify that the volume group was created successfully. Also type **pvs**. Notice that this command now shows the name of the physical volumes, with the names of the volume groups they are assigned to.
4. Type **lvcreate -n lvdata -l 50%FREE vgdata**. This creates an LVM logical volume with the name **lvdata**, which will use 50% of available disk space in the **vgdata** volume group.
5. Type **lvs** to verify that the volume was added successfully.
6. At this point, you are ready to create a file system on top of the logical volume. Type **mkfs.ext4 /dev/vgdata/lvdata** to create the file system.
7. Type **mkdir /files** to create a folder on which the volume can be mounted.
8. Add the following line to the bottom of **/etc/fstab**:

```
/dev/vgdata/lvdata /files ext4 defaults 0 0
```
9. Type **mount -a** to verify that the mount works and mount the file system.

Table 15-2 summarizes the relevant commands for creating logical volumes.

Table 15-2 LVM Management Essential Commands

**Key
Topic**

Command	Explanation
pvcreate	Creates physical volumes
pvs	Shows a summary of available physical volumes
pvdisplay	Shows a list of physical volumes and their properties
pvremove	Removes the physical volume signature from a block device
vgcreate	Creates volume groups
vgs	Shows a summary of available volume groups
vgdisplay	Shows a detailed list of volume groups and their properties
vgremove	Removes a volume group
lvcreate	Creates logical volumes
lvs	Shows a summary of all available logical volumes
lvdisplay	Shows a detailed list of available logical volumes and their properties
lvremove	Removes a logical volume

Resizing LVM Logical Volumes

One of the major benefits of using LVM is that LVM volumes are easy to resize, which is very useful if your file system is running out of available disk space. If the XFS file system is used, a volume can be increased, but not decreased, in size. Other file systems such as Ext4 support decreasing of the file system size also. Decreasing an Ext4 file system can be done offline only, which means that you need to unmount it before you can resize it. In this section, you learn how to increase the size of an LVM logical volume. To increase the size of a logical volume, you need to have disk space available in the volume group, so we address that first.

Resizing Volume Groups

The most important feature of LVM flexibility lies in the fact that it is so easy to resize the volume groups and the logical volumes that are using disk space from the volume groups. The **vgextend** command is used to add storage to a volume group, and the **vgreduce** command is used to take physical volumes out of a volume group

(which can lead to some additional complications). For the RHCSA exam, you need to know how to extend the available storage in volume groups. This procedure is relatively easy:

1. Make sure that a physical volume or device is available to be added to the volume group.
2. Use **vgextend** to extend the volume group. The new disk space will show immediately in the volume group.

After extending a volume group, you can use the **vgs** command to verify that a physical volume has been added to the volume group. In Example 15-8, you can see that the vgdata VG contains two physical volumes, as indicated in the #PV column.

Example 15-8 Verifying VG Resize Operations with **vgs**

```
[root@server1 ~]# vgs
```

VG	#PV	#LV	#SN	Attr	VSize	VFree
centos	1	2	0	wz--n-	<19.00g	0
vgdata	2	1	0	wz--n-	1020.00m	512.00m

Resizing Logical Volumes and File Systems

Like volume groups can be extended with the **vgextend** command, logical volumes can be extended with the **lvextend** command. This command has a very useful option **-r** to take care of extending the file systems on the logical volume at the same time; it is recommended to use this option and not the alternative approach that extends separately the logical volumes and the file systems on top of the logical volumes. Most file system resizing operations can be done online if the file system is getting bigger, without the need to unmount the file system.

To grow the logical volume size, use **lvextend** or **lvresize**, followed by the **-r** option to resize the file system used on it. Then specify the size you want the resized volume to be. The easiest and most intuitive way to do that is by using **-L** followed by a **+** sign and the amount of disk space you want to add, as in **lvresize -L +1G -r /dev/vgdata/lvdata**. An alternative way to resize the logical volume is by using the **-l** option. This option is followed either by the number of extents that are added to the logical volume or by the absolute or relative percentage of extents in the volume group that will be used. You can, for example, use the following commands to resize the logical volume:

- **lvresize -r -l 75%VG /dev/vgdata/lvdata** This resizes the logical volume so that it will take 75% of the total disk space in the volume group.



- **lvresize -r -l +75%VG /dev/vgdata/lvdata** This tries to add 75% of the total size of the volume group to the logical volume. (Notice the difference with the previous command.)
- **lvresize -r -l +75%FREE /dev/vgdata/lvdata** This adds 75% of all free disk space to the logical volume.
- **lvresize -r -l 75%FREE /dev/vgdata/lvdata** This resizes the logical volume to a total size that equals 75% of the amount of free disk space. (Notice the difference with the previous command.)

A logical extent is the logical building block used when creating logical volumes, and it maps to a physical extent, the size of which can be specified when creating a volume group. All resize operations need to match complete logical extents. You will sometimes notice that the resize size is rounded up or down to the logical extent size. You can also specify the number of logical extents that need to be added or removed directly by using the **-l** option with the **lvresize** command.

As you can see, resizing a logical volume has many options, and you need to take care to use the right options because it is easy to make a mistake! In Exercise 15-3, you learn how to resize logical volumes and the file systems used on top of them.

NOTE The size of an XFS file system cannot be decreased; it can only be increased. If you need a file system that can be shrunk in size, use Ext4, not XFS.

Exercise 15-3 Resizing Logical Volumes

In Exercises 15-1 and 15-2, you created a physical volume, volume group, and logical volume. In this exercise, you extend the size of the logical volume and the file system used on top of it.

1. Type **pvs** and **vgs** to show the current physical volume and volume group configuration.
2. Use **parted** to add another partition with a size of 1 GiB. Do not forget to flag this partition with the LVM partition type using **set lvm on**. I'll assume this new partition is **/dev/sdb2** for the rest of this exercise. Replace this name with the name used on your configuration if it is different.
3. Type **vgextend vgdata /dev/sdb2** to extend **vgdata** with the total size of the **/dev/sdb2** device.
4. Type **vgs** to verify that the available volume group size has increased.
5. Type **lvs** to verify the current size of the logical volume **lvdata**.

6. Type **df -h** to verify the current size of the file system on **lvdata**.
7. Type **lvextend -r -l +50%FREE /dev/vgdata/lvdata** to extend **lvdata** with 50% of all available disk space in the volume group.
8. Type **lvs** and **df -h** again to verify that the added disk space has become available.
9. Type **lvreduce -r -L -50M /dev/vgdata/lvdata**. This shrinks the **lvdata** volume with 50 MB. Notice that while doing this the volume is temporarily unmounted, which happens automatically. Also note that this step works only if you're using an Ext4 file system. (XFS cannot be shrunk.)

Configuring Stratis

In RHEL 8, Red Hat has added two advanced storage types: Stratis and VDO. Stratis is a so-called *volume-managing file system*, and it introduces advanced storage features that were not available in previous versions of RHEL:

- **Thin provisioning:** This feature enables a Stratis file to present itself to users as much bigger than it really is. This is useful in many environments, such as virtual desktops, where each user may see 20 GiB of available storage in total although a much lower amount is actually provisioned to each user.
- **Snapshots:** A Stratis snapshot allows users to take a “picture” of the current state of a file system. This snapshot makes it easy to revert to the previous state of a file system, rolling back any changes that have been made.
- **Cache tier:** Cache tier is a Ceph storage feature that ensures that data can be stored physically closer to the Ceph client, which makes data access faster.
- **Programmatic API:** The programmatic API ensures that storage can easily be configured and modified through API access. This is particularly interesting in cloud environments, where setting up storage directly from cloud-native applications is extremely useful.
- **Monitoring and repair:** Whereas old file systems need tools like **fsck** to verify the integrity of the file system, Stratis has built-in features to monitor the health of the file system and repair it if necessary.

Understanding Stratis Architecture

The lowest layer in the Stratis architecture is the pool. From a functional perspective, the Stratis pool is comparable to an LVM volume group. A *pool* represents all the available storage and consists of one or more storage devices, which in a

Stratis environment are referred to as *blockdev*. These block devices may not be thin provisioned at the underlying hardware level. Stratis creates a `/dev/stratis/poolname` directory for each pool.

From the Stratis pool, XFS file systems are created. Note that Stratis only works with XFS, and the XFS file system it uses is integrated with the Stratis volume. When creating a file system, no size is specified, and each file system can grow up to the size of all the available storage space in the pool. Stratis file systems are always thin provisioned. The thin volume automatically grows as more data is added to the file system.

Creating Stratis Storage

Creating Stratis volumes is a multistep process. This section provides a high-level overview, and then Exercise 15-4 in the following section guides you through the procedure. You start by creating a pool. Once the pool has been added, you can create file systems from it. Before you begin, make sure that the block devices you're going to use in Stratis have a minimal size of 1 GiB. Each Stratis file system occupies a minimum of 527 MiB of disk space, even if no data has been copied to the file system.

1. Install the Stratis software using **yum** by installing the **stratis-cli** and **stratisd** packages.
2. Start and enable the user-space daemon, using **systemctl enable --now stratisd**.
3. Once the daemon is running, use the **stratis pool create** command to create the pool that you want to work with. For instance, use **stratis pool create mypool /dev/sdd** to create a pool that is based on the block device `/dev/sdd`. You can add additional block devices later, using **stratis pool add-data poolname blockdevname**, as in **stratis pool add-data mypool /dev/sde**.
4. Once you have created the pool, add a file system using **stratis fs create poolname fsname**.
5. To verify that all was created successfully, use the **stratis fs list** command.
6. After creating the file system, you can mount it. To mount a Stratis file system through `/etc/fstab`, you *must* use the **UUID**; using the device name is not supported. Also, when mounting the Stratis volume through `/etc/fstab`, include the mount option **s-systemd.requires=stratisd.service** to ensure that the Systemd waits to activate this device until the **stratisd** service is loaded.

Managing Stratis

After creating the Stratis file system, you can perform several different management tasks. To start with, you can dynamically extend the pool, using **stratis pool add-data**. Also, you need to monitor Stratis volumes using Stratis-specific tools, as the traditional Linux tools cannot handle the thin-provisioned volumes. The following commands are available:

- **stratis blockdev**: Shows information about all block devices that are used for Stratis.
- **stratis pool**: Gives information about Stratis pools. Note in particular the Physical Used parameter, which should not come too close to the Physical Size parameter.
- **stratis filesystem**: Enables you to monitor individual file systems.

Another Stratis feature that you'll have to manage is the snapshot. A snapshot contains the state of the file system at the moment the snapshot was created. After creation, the snapshot can be modified. It's also good to know that the snapshot and its origin are not linked, which allows the snapshot to live longer than the file system it was created from. This is fundamentally different from, for instance, LVM snapshots, which cannot stay alive if the volume they are linked to is removed.

In Exercise 15-4, you set up an environment with Stratis volumes.

Exercise 15-4 Managing Stratis Volumes

You need one dedicated disk with a size of 5 GiB to perform the steps in this exercise. In this exercise, the disk name `/dev/sdd` is used as an example. Replace this name with the disk device name that is presented on your hardware.

1. Type **yum install stratisd stratis-cli** to install all the required packages.
2. Type **systemctl enable --now stratisd** to enable the Stratis daemon.
3. Type **stratis pool create mypool /dev/sdd** to add the entire disk `/dev/sdd` to the storage pool.
4. Type **stratis pool list** to verify successful creation of the pool.
5. Type **stratis fs create mypool stratis1** to create the first Stratis file system. Note that you don't have to specify a file system size.
6. Type **stratis fs list** to verify the creation of the file system.
7. Type **mkdir /stratis1** to create a mount point for the Stratis file system.
8. Type **blkid** to find the Stratis volume UUID.

9. Add the following line to `/etc/fstab` to enable the volume to be mounted automatically. Make sure to use the UUID name that is used on your system.

```
UUID=xxx /stratis1. xfs. defaults,x-systemd.requires=stratisd.  
service 0 0
```
10. Type **mount -a** to mount the Stratis volume. Use the **mount** command to verify this worked successfully.
11. Type **cp /etc/[a-f]* /stratis1** to copy some files to the Stratis volume.
12. Type **stratis filesystem snapshot mypool stratis1 stratis1-snap** to create a snapshot of the volume you just created. Note that this command may take up to a minute to complete.
13. Type **stratis filesystem list** to get statistics about current file system usage.
14. Type **rm -f /stratis1/a*** to remove all files that have a name starting with *a*.
15. Type **mount /dev/stratis/mypool/stratis1-snap /mnt** and verify that the files whose names start with *a* are still available in the `/mnt` directory.
16. Reboot your server. After reboot, verify that the Stratis volume is still automatically mounted.

Configuring VDO

In the previous section you learned about Stratis. In this section you'll learn about Virtual Data Optimizer (VDO), a solution that was developed to offer data deduplication features.

Understanding VDO

Virtual Data Optimizer (VDO) is another advanced storage solution offered in RHEL 8. VDO was developed to reduce disk space usage on block devices by applying deduplication features. VDO creates volumes, implementing deduplication on top of any type of existing block device. On top of the VDO device, you would either create a file system, or you would use it as a physical volume in an LVM setup.

The main purpose of VDO is data deduplication. It deduplicates data by using three common technologies:

- Zero-block elimination filters out data blocks that contain only zeros.
- Deduplication reduces redundant data blocks.
- Compression occurs when the `kvdo` kernel module compresses data blocks.

VDO is particularly useful in specific environments. The typical use cases are host platforms for containers and virtual machines, where the deduplication engine ensures much more efficient disk storage use. Commonly, a logical size of up to ten times the physical size can be used for working in this type of environment.

Another typical environment where VDO does well is in cloud block storage. If, for instance, VDO is used as the underlying storage layer for Ceph storage, a logical size of up to three times the physical size can be used.

Setting Up VDO

To configure VDO, you first need to make sure that the underlying block devices have a minimal size of 4 GiB. Next, you need to install the **vdo** and **kmod-vdo** packages. These ensure that the required kernel modules are available.

Next, create the device by using the **vdo create** command. While creating the device, you can use the **--vdiLogicalSize=** option to specify the logical size of the volume. For instance, use **vdo create --name=vdo1 --device=/dev/sde --vdoLogicalSize=1T** to create a volume with a logical size of 1 TiB.

After creating the device, you can put an XFS file system on top of it, using **mkfs.xfs -K /dev/mapper/vdo1**. The **-K** option with **mkfs.xfs** prevents unused blocks in the file system from being discarded immediately, which makes the command significantly faster.

Once the file system has been created, you can add a line to **/etc/fstab** to automatically mount it. Make sure to include the mount options **x-systemd.requires=vdo.service,discard**, which triggers the **vdo systemd** service to be loaded. To verify that all is working well, you can now use the command **vdstats --human-readable**.

TIP Memorizing the options to use to make a VDO mount persistent may be challenging. Fortunately, in **/usr/share/doc/vdo/examples/systemd** there is an example Systemd mount file, shown in Example 15-9. Just modify the settings of the mount file to your requirements, and copy it to **/etc/systemd/system/yourmountpoint.mount**. Don't forget to enable it, and it will automatically be mounted at the next reboot.

Example 15-9 Sample Systemd Mount Unit File for VDO

```
[Unit]
Description = Mount filesystem that lives on VDO
name = VDO.mount
Requires = vdo.service systemd-remount-fs.service
After = multi-user.target
Conflicts = umount.target
```

```
[Mount]
What = /dev/mapper/my_vdo
Where = /VDO
Type = xfs
Options = discard

[Install]
WantedBy = multi-user.target
[root@localhost systemd]# pwd
/usr/share/doc/vdo/examples/systemd
```

Exercise 15-5 guides you through the steps to create a thin-provisioned VDO volume.

Exercise 15-5 Managing VDO Storage

To do this exercise, you need to have access to a block device with a size of at least 5 GiB. This exercise uses `/dev/sde` as an example block device name. Make sure to replace this name with the actual device name that is used in your configuration.

1. Type **lsblk** to verify the availability of the `/dev/sde` block device.
2. Type **yum install -y vdo kmod-kvdo** to install the required packages.
3. Type **vdo create --name=vdo1 --device=/dev/sde --vdoLogicalSize=1T** to create the VDO device with a logical size of 1 TiB.
4. Use **mkfs.xfs -K /dev/mapper/vdo1** to put an XFS file system on top of the device.
5. Type **mkdir /vdo1** to create a mount point where the VDO device can be automatically mounted.
6. Copy the example Systemd mount file for VDO to `/etc/systemd/system`, using **cp /usr/share/doc/vdo/examples/systemd/VDO.mount.example /etc/systemd/system/vdo1.mount**.
7. Edit the file to include the following lines:

```
what = /dev/mapper/vdo1
where = /vdo1
```
8. Type **systemctl enable --now vdo1.mount** to mount the new block device.
9. Type **vdo stats --human-readable** to monitor current statistics about the VDO device.
10. Type **df -h** to verify the logical size of 1 GiB.
11. Reboot your server to ensure that after reboot the VDO device is correctly mounted.

Summary

In this chapter, you learned how to work with advanced storage on RHEL 8. First, you read how LVM is used to bring flexibility to storage. By using LVM, you get the advantages of volumes that can be resized easily and multi-device logical volumes. Next, you were introduced to Stratis, the volume-managing file system. Stratis brings next-generation storage features to RHEL 8, and by default creates thin-provisioned file systems. Last, you learned how to use VDO, the Virtual Data Optimizer, to implement data deduplication in specific scenarios.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 15-3 lists a reference of these key topics and the page number on which each is found.



Table 15-3 Key Topics for Chapter 15

Key Topic Element	Description	Page
Figure 15-1	LVM architecture overview	343
Table 15-2	LVM management essential commands	353
List	LVM lvresize commands	354

Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

physical volume (PV), volume group (VG), logical volume (LV), device mapper, physical extent, logical extent, snapshot, Stratis, VDO, deduplication, thin allocation

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. Which partition type is used on a GUID partition that needs to be used in LVM?
2. Which command enables you to create a volume group with the name `vgroup` that contains the physical device `/dev/sdb3` and uses a physical extent size of 4 MiB?
3. Which command shows a short summary of the physical volumes on your system as well as the volume group to which these belong?
4. What do you need to do to add an entire hard disk `/dev/sdd` to the volume group `vgroup`?
5. Which command enables you to create a logical volume `lvvol1` with a size of 6 MiB?
6. Which command enables you to add 100 MB to the logical volume `lvvol1`, assuming that the disk space is available in the volume group?
7. When working with Stratis, what line would you add to `/etc/fstab` to mount the Stratis volume?
8. Which command do you use to create a Stratis pool that is based on the block device `/dev/sdd`?
9. Which command do you use to create a VDO device with a virtual size of 1 TiB?
10. Which option do you need with `mkfs.xfs` to successfully create an XFS file system on top of a VDO thin-provisioned volume?

End-of-Chapter Labs

To complete the following end-of-chapter labs, you need a dedicated disk device. Either use a USB thumb drive or add a new virtual disk to your virtual environment before starting.

Lab 15.1

1. Create a 500-MB logical volume named **lvgroup**. Format it with the XFS file system and mount it persistently on `/groups`. Reboot your server to verify that the mount works.
2. After rebooting, add another 250 MB to the `lvgroup` volume that you just created. Verify that the file system resizes as well while resizing the volume.
3. Verify that the volume extension was successful.

Lab 15.2

1. Create a Stratis pool with a size of 5 GiB. In this pool, create two Stratis file systems and ensure that they are automatically mounted.
2. Add an additional block device to the Stratis pool and verify that the size of the pool was successfully extended.

Lab 15.3

1. Create a VDO volume with a virtual size of 1 TiB.
2. Ensure the VDO volume is automatically mounted by using the example Systemd unit file.
3. Reboot your server to verify that everything still works.

This page intentionally left blank



The following topics are covered in this chapter:

- Understanding the Role of the Linux Kernel
- Working with Kernel Modules
- Upgrading the Linux Kernel

The Linux kernel is the heart of the Linux operating system. It takes care of many things, including hardware management. In this chapter, you learn all you need to know about the Linux kernel from an RHCSA perspective. In fact, you even learn a bit more. This chapter includes information about topics that are not on the current list of RHCSA objectives. I think it is good to know about these topics anyway. Any serious Linux administrator should be able to deal with issues related to the topics discussed in this chapter.

Basic Kernel Management

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 16-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 16-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Understanding the Role of the Linux Kernel	1–4
Working with Kernel Modules	5–9
Upgrading the Linux Kernel	10

1. What causes a tainted kernel?
 - a. A kernel driver that is not available as an open source driver
 - b. A driver that was developed for a different operating system but has been ported to Linux
 - c. A driver that has failed
 - d. An unsupported driver
2. What is the name of the command that shows kernel events since booting?
 - a. `logger`
 - b. `dmesg`
 - c. `klogd`
 - d. `journald`

3. Which command enables you to find the actual version of the kernel that is used?
 - a. `uname -r`
 - b. `uname -v`
 - c. `procinfo -k`
 - d. `procinfo -l`
4. Which command shows the current version of RHEL you are using?
 - a. `uname -r`
 - b. `cat /proc/rhel-version`
 - c. `cat /etc/redhat-release`
 - d. `uname -k`
5. What is the name of the process that helps the kernel to initialize hardware devices properly?
 - a. `systemd-udev`
 - b. `hwinit`
 - c. `udev`
 - d. `udev`
6. Where does your system find the default rules that are used for initializing new hardware devices?
 - a. `/etc/udev/rules.d`
 - b. `/usr/lib/udev/rules.d`
 - c. `/usr/lib/udev.d/rules`
 - d. `/etc/udev.d/rules`
7. Which command should you use to unload a kernel module, including all of its dependencies?
 - a. `rmmod`
 - b. `insmod -r`
 - c. `modprobe -r`
 - d. `modprobe`

8. Which command enables you to see whether the appropriate kernel modules have been loaded for hardware in your server?
 - a. **lsmod**
 - b. **modprobe -l**
 - c. **lspci -k**
 - d. **lspci**
9. Where do you specify a kernel module parameter to make it persistent?
 - a. `/etc/modules.conf`
 - b. `/etc/modprobe.conf`
 - c. `/etc/modprobe.d/somefilename`
 - d. `/usr/lib/modprobe.d/somefilename`
10. Which statements about updating the kernel are *not* true?
 - a. The **yum update kernel** command will install a new kernel and not update it.
 - b. The **yum install kernel** command will install a new kernel and keep the old kernel.
 - c. The kernel package should be set as a **yum**-protected package to ensure that after an update the old kernel is still available.
 - d. After you have installed a new kernel version, you must run the **grub2-mkconfig** command to modify the GRUB 2 boot menu so that it shows the old kernel and the newly installed kernel.

Foundation Topics

Understanding the Role of the Linux Kernel

The Linux kernel is the heart of the operating system. It is the layer between the user who works with Linux from a shell environment and the hardware that is available in the computer on which the user is working. The kernel manages the I/O instructions it receives from the software and translates them into the processing instructions that are executed by the central processing unit and other hardware in the computer. The kernel also takes care of handling essential operating system tasks. One example of such a task is the scheduler that makes sure any processes that are started on the operating system are handled by the CPU.

Understanding the Use of Kernel Threads and Drivers

The operating system tasks that are performed by the kernel are implemented by different kernel threads. Kernel threads are easily recognized with a command like **ps aux**. The kernel thread names are listed between square brackets (see Example 16-1).

Example 16-1 Listing Kernel Threads with **ps aux**

```
[root@server1 ~]# ps aux | head -n 20
```

USER	PID	%CPU	%MEM	VSZ	RSS	TTY	STAT	START	TIME	COMMAND
root	1	1.8	0.6	52980	6812	?	Ss	11:44	0:02	/usr/lib/systemd/systemd --switched-root --system --deserialize 23
root	2	0.0	0.0	0	0	?	S	11:44	0:00	[kthreadd]
root	3	0.0	0.0	0	0	?	S	11:44	0:00	[ksoftirqd/0]
root	4	0.0	0.0	0	0	?	S	11:44	0:00	[kworker/0:0]
root	5	0.0	0.0	0	0	?	S<	11:44	0:00	[kworker/0:0H]
root	6	0.0	0.0	0	0	?	S	11:44	0:00	[kworker/u128:0]
root	7	0.1	0.0	0	0	?	S	11:44	0:00	[migration/0]
root	8	0.0	0.0	0	0	?	S	11:44	0:00	[rcu_bh]
root	9	0.0	0.0	0	0	?	S	11:44	0:00	[rcuob/0]
root	10	0.0	0.0	0	0	?	S	11:44	0:00	[rcuob/1]
root	11	0.0	0.0	0	0	?	S	11:44	0:00	[rcuob/2]
root	12	0.0	0.0	0	0	?	S	11:44	0:00	[rcuob/3]
root	13	0.0	0.0	0	0	?	S	11:44	0:00	[rcuob/4]
root	14	0.0	0.0	0	0	?	S	11:44	0:00	[rcuob/5]
root	15	0.0	0.0	0	0	?	S	11:44	0:00	[rcuob/6]
root	16	0.0	0.0	0	0	?	S	11:44	0:00	[rcuob/7]
root	17	0.0	0.0	0	0	?	S	11:44	0:00	[rcuob/8]
root	18	0.0	0.0	0	0	?	S	11:44	0:00	[rcuob/9]
root	19	0.0	0.0	0	0	?	S	11:44	0:00	[rcuob/10]

Another important task of the Linux kernel is hardware initialization. To make sure that this hardware can be used, the Linux kernel uses drivers. Every piece of hardware contains specific features, and to use these features, a driver must be loaded. The Linux kernel is modular, and drivers are loaded as kernel modules, which you read about more later in this chapter.

In some cases, the availability of drivers is an issue because hardware manufacturers are not always willing to provide open source drivers that can be integrated well with the Linux kernel. That can result in a driver that does not provide all the functionality that is provided by the hardware.

If a manufacturer is not willing to provide open source drivers, an alternative is to work with closed source drivers. Although these make it possible to use the hardware in Linux, the solution is not ideal. Because a driver performs privileged instructions within the kernel space, a badly functioning driver may crash the entire kernel. If this happens with an open source driver, the Linux kernel community can help debug the problem and make sure that the issue is fixed. If it happens with a closed source driver, the Linux kernel community cannot do anything. But, a proprietary driver may provide access to features that are not provided by its open source equivalent.

To make it easy to see whether a kernel is using closed source drivers, the concept of the tainted kernel is used. A *tainted kernel* is a kernel that contains closed source drivers. The concept of tainted kernels helps in troubleshooting drivers. If your RHEL 8 kernel appears to be tainted, Red Hat support can identify it as a tainted kernel and recognize which driver is tainting it. To fix the problem, Red Hat might ask you to take out the driver that is making it a tainted kernel.

Analyzing What the Kernel Is Doing

To help analyze what the kernel is doing, some tools are provided by the Linux operating systems:

- The **dmesg** utility
- The **/proc** file system
- The **uname** utility

**Key
Topic**

The first utility to consider if you require detailed information about the kernel activity is **dmesg**. This utility shows the contents of the kernel ring buffer, an area of memory where the Linux kernel keeps its recent log messages. An alternative method to get access to the same information in the kernel ring buffer is the **journalctl --dmesg** command, which is equivalent to **journalctl -k**. In Example 16-2, you can see a part of the result of the **dmesg** command.

Example 16-2 Analyzing Kernel Activity Using **dmesg**

```
[ 8.153928] sd 0:0:0:0: Attached scsi generic sg0 type 0
[ 8.154289] sd 0:0:1:0: Attached scsi generic sg1 type 0
[ 8.154330] sd 0:0:2:0: Attached scsi generic sg2 type 0
[ 8.154360] sd 0:0:3:0: Attached scsi generic sg3 type 0
[ 8.154421] sr 4:0:0:0: Attached scsi generic sg4 type 5
[ 8.729016] ip_tables: (C) 2000-2006 Netfilter Core Team
[ 8.850599] nf_conntrack version 0.5.0 (7897 buckets, 31588 max)
[ 8.939613] ip6_tables: (C) 2000-2006 Netfilter Core Team
[ 9.160092] Ebtables v2.0 registered
[ 9.203710] Bridge firewalling registered
[ 9.586603] IPv6: ADDRCONF(NETDEV_UP): eno16777736: link is not
ready
[ 9.587520] e1000: eno16777736 NIC Link is Up 1000 Mbps Full
Duplex, Flow Control: None
[ 9.589066] IPv6: ADDRCONF(NETDEV_CHANGE): eno16777736: link becomes
ready
[ 10.689365] Rounding down aligned max_sectors from 4294967295 to
4294967288
[ 5158.470480] Adjusting tsc more than 11% (6940512 vs 6913395)
[21766.132181] e1000: eno16777736 NIC Link is Down
[21770.391597] e1000: eno16777736 NIC Link is Up 1000 Mbps Full
Duplex, Flow Control: None
[21780.434547] e1000: eno16777736 NIC Link is Down
```

In the **dmesg** output, all kernel-related messages are shown. Each message starts with a time indicator that shows at which specific second the event was logged. This time indicator is relative to the start of the kernel, which allows you to see exactly how many seconds have passed between the start of the kernel and a particular event. (Notice that the **journalctl -k / --dmesg** commands show clock time, instead of time that is relative to the start of the kernel.) This time indicator gives a clear indication of what has been happening and at which time it has happened.

Another valuable source of information is the **/proc** file system. The **/proc** file system is an interface to the Linux kernel, and it contains files with detailed status information about what is happening on your server. Many of the performance-related tools mine the **/proc** file system for more information.

As an administrator, you will find that some of the files in **/proc** are very readable and contain status information about the CPU, memory, mounts, and more. Take a look, for instance, at **/proc/meminfo**, which gives detailed information about each memory segment and what exactly is happening in these memory segments.

A last useful source of information is the **uname** command. This command gives different kinds of information about your operating system. Type, for instance, **uname -a** for an overview of all relevant parameters of **uname -r** to see which kernel version currently is used. This information also shows when using the **hostnamectl status** command, which shows useful additional information as well (see Example 16-3).

TIP On some occasions, you might need to know specific information about the RHEL version you are using. To get that information, run the **/etc/redhat-release** command and review its output; it will tell you which Red Hat version you are using and which update level is applied.

Example 16-3 Getting More Information About the System

```
[root@server1 ~]# hostnamectl status
  Static hostname: server1.example.com
        Icon name: computer-vm
        Chassis: vm
  Machine ID: 5aa095b495ed458d934c54a88078c165
  Boot ID: b7273a66ba254358b566639e329c350d
  Virtualization: vmware
  Operating System: Red Hat Enterprise Linux 8.0 (Ootpa)
    CPE OS Name: cpe:/o:redhat:enterprise_linux:8.0:GA
        Kernel: Linux 4.18.0-80.el8.x86_64
  Architecture: x86-64
```

Working with Kernel Modules

In the old days of Linux, kernels had to be compiled to include all drivers that were required to support computer hardware. Other specific functionality needed to be compiled into the kernel as well. Since the release of Linux kernel 2.0 in the mid-1990s, kernels are no longer compiled but modular. A *modular kernel* consists of a relatively small core kernel and provides driver support through modules that are loaded when required. Modular kernels are very efficient, as they include only those modules that really are needed.

TIP A kernel module implements specific kernel functionality. Kernel modules are used to load drivers that allow proper communications with hardware devices, but are not limited to loading hardware drivers alone. For example, file system support is loaded as modules. Other kernel features can be loaded as modules as well.

Understanding Hardware Initialization

The loading of drivers is an automated process that roughly goes like this:

1. During boot, the kernel probes available hardware.
2. Upon detection of a hardware component, the **systemd-udev** process takes care of loading the appropriate driver and making the hardware device available.
3. To decide how the devices are initialized, **systemd-udev** reads rules files in `/usr/lib/udev/rules.d`. These are system-provided rules files that should not be modified.
4. After processing the system-provided udev rules files, **systemd-udev** goes to the `/etc/udev/rules.d` directory to read any custom rules if these are available.
5. As a result, required kernel modules are loaded automatically, and status about the kernel modules and associated hardware is written to the `sysfs` file system, which is mounted on the `/sys` directory. The Linux kernel uses this pseudo file system to track hardware-related settings.

The **systemd-udev** process is not a one-time-only process; it continuously monitors plugging and unplugging of new hardware devices. To get an impression of how this works, as root you can type the command **udevadm monitor**. This lists all events that are processed while activating new hardware devices. For instance, if you plug in a USB device while this command is active, you can see exactly what's happening. Press Ctrl-C to close the **udevadm monitor** output.

Example 16-4 shows output of the **udevadm monitor** command. In this command, you can see how features that are offered by the hardware are discovered automatically by the kernel and **systemd-udev** working together. Each phase of the hardware probing is concluded by the creation of a file in the `/sys` file system. Once the hardware has been fully initialized, you can also see that some kernel modules are loaded.

NOTE Although useful to know, hardware initialization is not included in the current RHCSA objectives.

Example 16-4 Output of the **udevadm monitor** Command

```
[root@server2 ~]# udevadm monitor
monitor will print the received events for:
UDEV - the event which udev sends out after rule processing
KERNEL - the kernel uevent

KERNEL[132406.831270] add /devices/pci0000:00/0000:00:11.0/0000:02:
04.0/usb1/l-1 (usb)
KERNEL[132406.974110] add /devices/pci0000:00/0000:00:11.0/0000:02:
04.0/usb1/l-1/l-1:1.0 (usb)
UDEV [132406.988182] add /devices/pci0000:00/0000:00:11.0/0000:02:
04.0/usb1/l-1 (usb)
KERNEL[132406.999249] add /module/usb_storage (module)
UDEV [132407.001203] add /module/usb_storage (module)
KERNEL[132407.002559] add /devices/pci0000:00/0000:00:11.0/0000:02:
04.0/usb1/l-1/l-1:1.0/host33 (scsi)
UDEV [132407.002575] add /devices/pci0000:00/0000:00:11.0/0000:02:
04.0/usb1/l-1/l-1:1.0 (usb)
KERNEL[132407.002583] add /devices/pci0000:00/0000:00:11.0/0000:02:
04.0/usb1/l-1/l-1:1.0/host33/scsi_host/host33 (scsi_host)
KERNEL[132407.002590] add /bus/usb/drivers/usb-storage (drivers)
UDEV [132407.004479] add /bus/usb/drivers/usb-storage (drivers)
UDEV [132407.005798] add /devices/pci0000:00/0000:00:11.0/0000:02:
04.0/usb1/l-1/l-1:1.0/host33 (scsi)
UDEV [132407.007385] add /devices/pci0000:00/0000:00:11.0/0000:02:
04.0/usb1/l-1/l-1:1.0/host33/scsi_host/host33 (scsi_host)
KERNEL[132408.008331] add /devices/pci0000:00/0000:00:11.0/0000:02:
04.0/usb1/l-1/l-1:1.0/host33/target33:0:0 (scsi)
KERNEL[132408.008355] add /devices/pci0000:00/0000:00:11.0/0000:02:
04.0/usb1/l-1/l-1:1.0/host33/target33:0:0/33:0:0:0 (scsi)
...
KERNEL[132409.381930] add /module/fat (module)
KERNEL[132409.381951] add /kernel/slab/fat_cache (slab)
KERNEL[132409.381958] add /kernel/slab/fat_inode_cache (slab)
KERNEL[132409.381964] add /module/vfat (module)
UDEV [132409.385090] add /module/fat (module)
UDEV [132409.385107] add /kernel/slab/fat_cache (slab)
UDEV [132409.385113] add /kernel/slab/fat_inode_cache (slab)
UDEV [132409.386110] add /module/vfat (module)
```

Managing Kernel Modules

Linux kernel modules normally are loaded automatically for the devices that need them, but you will on rare occasions have to load the appropriate kernel modules manually. A few commands are used for manual management of kernel modules. Table 16-2 provides an overview.

An alternative method of loading kernel modules is through the `/etc/modules-load.d` directory. In this directory, you can create files to load modules automatically that are not loaded by the `systemd-udev` method already. For default modules that should always be loaded, this directory has a counterpart in `/usr/lib/modules-load.d`.

Key Topic

Table 16-2 Linux Kernel Module Management Overview

Command	Use
lsmod	Lists currently loaded kernel modules
modinfo	Displays information about kernel modules
modprobe	Loads kernel modules, including all of their dependencies
modprobe -r	Unloads kernel modules, considering kernel module dependencies

The first command to use when working with kernel modules is **lsmod**. This command lists all kernel modules that currently are used, including the modules by which this specific module is used. Example 16-5 shows the output of the first ten lines of the **lsmod** command.

Example 16-5 Listing Loaded Modules with **lsmod**

```
[root@server1 ~]# lsmod | head
Module                Size      Used by
nls_utf8              16384      1
isofs                 45056      1
fuse                 126976      3
rfcomm                90112      6
xt_CHECKSUM           16384      1
ipt_MASQUERADE        16384      1
xt_conntrack          16384      1
ipt_REJECT            16384      1
nft_counter           16384     16
```

TIP Many Linux commands show their output in different columns, and it is not always clear which column is used to show which kind of information. Most of these commands have a header line on the first line of command output. So, if in the output of any command you are not sure what you are seeing, pipe the output of the command through **head** to see whether there is a header file, or pipe the command output to **less**, which allows you to page up to the first line of command output easily.

If you want to have more information about a specific kernel module, you can use the **modinfo** command. This gives complete information about the specific kernel modules, including two interesting sections: the alias and the parms. A module alias is another name that can also be used to address the module. The parms lines refer to parameters that can be set while loading the module. (In the section “Managing Kernel Module Parameters” later in this chapter, you learn how to work with kernel module parameters.) Example 16-6 shows partial output of the **modinfo e1000** command.

Example 16-6 Showing Module Information with **modinfo**

```
[root@server1 ~]# modinfo e1000
filename:          /lib/modules/4.18.0-80.el8.x86_64/kernel/drivers/net/
                  ethernet/intel/e1000/e1000.ko.xz
version:           7.3.21-k8-NAPI
license:           GPL
description:       Intel(R) PRO/1000 Network Driver
author:            Intel Corporation, <linux.nics@intel.com>
rhelversion:       8.0
srcversion:        A2C44EC5D0B865EE9C972D5
alias:             pci:v00008086d00002E6Esv*sd*bc*sc*i*
...
depends:
intree:            Y
name:              e1000
vermagic:          4.18.0-80.el8.x86_64 SMP mod_unload modversions
sig_id:            PKCS#7
signer:            Red Hat Enterprise Linux kernel signing key
sig_key:           20:16:97:CB:B1:7E:D5:A0:A7:3C:0B:40:E2:54:80:2D:83:F1:
                  72:5A
sig_hashalgo:      sha256
signature:         97:3A:58:F6:4E:B7:F1:CE:44:50:65:5E:84:41:D8:A7:A3:
                  A2:2D...
parm:              TxDescriptors:Number of transmit descriptors (array
                  of int)
```

```

parm:      RxDescriptors: Number of receive descriptors (array
              of int)
parm:      Speed: Speed setting (array of int)
parm:      Duplex: Duplex setting (array of int)
parm:      AutoNeg: Advertised auto-negotiation setting (array
              of int)
parm:      FlowControl: Flow Control setting (array of int)
parm:      XsumRX: Disable or enable Receive Checksum offload
              (array of int)
parm:      TxIntDelay: Transmit Interrupt Delay (array of int)
parm:      TxAbsIntDelay: Transmit Absolute Interrupt Delay
              (array of int)
parm:      RxIntDelay: Receive Interrupt Delay (array of int)
parm:      RxAbsIntDelay: Receive Absolute Interrupt Delay
              (array of int)
parm:      InterruptThrottleRate: Interrupt Throttling Rate
              (array of int)
parm:      SmartPowerDownEnable: Enable PHY smart power down
              (array of int)
parm:      copybreak: Maximum size of packet that is copied to
              a new buffer on receive (uint)
parm:      debug: Debug level (0=none,...,16=all) (int)

```

To manually load and unload modules, you can use the **modprobe** and **modprobe -r** commands. On earlier Linux versions, you may have used the **insmod** and **rmmod** commands. These should no longer be used because they do not load kernel module dependencies. In Exercise 16-1, you learn how to manage kernel modules using these commands.

Exercise 16-1 Managing Kernel Modules from the Command Line

1. Open a root shell and type **lsmod | less**. This shows all kernel modules currently loaded.
2. Type **modprobe ext4** to load the ext4 kernel module.
3. Verify that the module is loaded by using the **lsmod | grep ext** command. You can see that the module is loaded, as well as some of its dependencies.
4. Type **modinfo ext4** to get information about the ext4 kernel module. Notice that it does not have any parameters.
5. Type **modprobe -r ext4** to unload the ext4 kernel module again.
6. Type **modprobe -r xfs** to try to unload the xfs kernel module. Notice that you get an error message because the kernel module currently is in use.

Checking Driver Availability for Hardware Devices

On modern Linux servers, many hardware devices are supported. On occasion, you might find that some devices are not supported properly because their modules are not currently loaded. The best way to find out whether this is the case for your hardware is by using the **lspci** command. If used without arguments, it shows all hardware devices that have been detected on the PCI bus. A very useful argument is **-k**, which lists all kernel modules that are used for the PCI devices that were detected. Example 16-7 shows sample output of the **lspci -k** command.

Example 16-7 Checking Kernel Module Availability

```
[root@server1 ~]# lspci -k | head
00:00.0 Host bridge: Intel Corporation 440BX/ZX/DX - 82443BX/ZX/DX
    Host bridge (rev 01)
        Subsystem: VMware Virtual Machine Chipset
        Kernel driver in use: agpgart-intel
00:01.0 PCI bridge: Intel Corporation 440BX/ZX/DX - 82443BX/ZX/
    DX AGP bridge (rev 01)
00:07.0 ISA bridge: Intel Corporation 82371AB/EB/MB PIIX4 ISA (rev 08)
    Subsystem: VMware Virtual Machine Chipset
00:07.1 IDE interface: Intel Corporation 82371AB/EB/MB PIIX4 IDE
    (rev 01)
    Subsystem: VMware Virtual Machine Chipset
    Kernel driver in use: ata_piix
    Kernel modules: ata_piix, ata_generic
00:07.3 Bridge: Intel Corporation 82371AB/EB/MB PIIX4 ACPI (rev 08)
    Subsystem: VMware Virtual Machine Chipset
    Kernel modules: i2c_piix4
00:07.7 System peripheral: VMware Virtual Machine Communication
    Interface (rev 10)
    Subsystem: VMware Virtual Machine Communication Interface
    Kernel driver in use: vmw_vmci
    Kernel modules: vmw_vmci
00:0f.0 VGA compatible controller: VMware SVGA II Adapter
    Subsystem: VMware SVGA II Adapter
    Kernel driver in use: vmwgfx
    Kernel modules: vmwgfx
00:10.0 SCSI storage controller: LSI Logic / Symbios Logic 53c1030
    PCI-X Fusion-MPT Dual Ultra320 SCSI (rev 01)
    Subsystem: VMware LSI Logic Parallel SCSI Controller
    Kernel driver in use: mptspi
    Kernel modules: mptspi
00:11.0 PCI bridge: VMware PCI bridge (rev 02)
00:15.0 PCI bridge: VMware PCI Express Root Port (rev 01)
    Kernel driver in use: pcieport
```

If you discover that PCI devices were found for which no kernel modules could be loaded, you are probably dealing with a device that is not supported. You can try to find a closed source kernel module, but you should realize that that might endanger the stability of your kernel. A much better approach is to check with your hardware vendor that Linux is fully supported before you purchase specific hardware.

Managing Kernel Module Parameters

Occasionally, you might want to load kernel modules with specific parameters. To do so, you first need to find out which parameter you want to use. If you have found the parameter you want to use, you can load it manually, specifying the name of the parameter followed by the value that you want to assign. To make this an automated procedure, you can create a file in the `/etc/modprobe.d` directory, where the module is loaded including the parameter you want to be loaded. In Exercise 16-2 you see how to do this using the `cdrom` kernel module.

Exercise 16-2 Loading Kernel Modules with Parameters

1. Type **`lsmod | grep cdrom`**. If you have used the optical drive in your computer, this module should be loaded, and it should indicate that it is used by the `sr_mod` module.
2. Type **`modprobe -r cdrom`**. This will not work because the module is in use by the `sr_mod` module.
3. Type **`modprobe -r sr_mod; modprobe -r cdrom`**. This should unload both modules, but it will most likely fail. (It won't fail if currently no optical device is mounted.)
4. Type **`umount /dev/sr0`** to unmount the mounted `cdrom` file system and perform step 3 of this exercise again.
5. Type **`modinfo cdrom`**. This shows information about the `cdrom` module, including the parameters that it supports. One of these is the **`debug`** parameter, which supports a Boolean as its value.
6. Type **`modprobe cdrom debug=1`**. This loads the `cdrom` module with the **`debug`** parameter set to on.
7. Type **`dmesg`**. For some kernel modules, load information is written to the kernel ring buffer, which can be displayed using the **`dmesg`** command. Unfortunately, this is not the case for the `cdrom` kernel module.
8. Create a file with the name `/etc/modprobe.d/cdrom.conf` and give it the following contents:

```
options cdrom debug=1
```

This enables the parameter every time the `cdrom` kernel module loads.

Upgrading the Linux Kernel

From time to time, you need to upgrade the Linux kernel. When you upgrade the Linux kernel, a new version of the kernel is installed and used as the default kernel. The old version of the kernel file will still be available, though. This ensures that your computer can still boot if in the new kernel nonsupported functionality is included. To install a new version of the kernel, you can use the command **yum upgrade kernel**. The **yum install kernel** command also works. Both commands install the new kernel beside the old kernel.

The kernel files for the last four kernels that you have installed on your server will be kept in the `/boot` directory. The GRUB 2 boot loader automatically picks up all kernels that it finds in this directory. This allows you to select an older kernel version while booting, which is useful if the newly installed kernel doesn't boot correctly.

Summary

In this chapter, you learned how to work with the Linux kernel. You learned that the Linux kernel is modular and how working with kernel modules is important. You also learned how to manage kernel modules and how kernel modules are managed automatically while working with new hardware.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 16-3 lists a reference of these key topics and the page number on which each is found.

Table 16-3 Key Topics for Chapter 16

Key Topic Element	Description	Page
List	Overview of kernel-related tools	371
Table 16-2	Kernel module management overview	376

Key
Topic

Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

kernel, module, **dmesg**, udev, sysfs, proc, tainted kernel

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. Which command shows the current version of the kernel that is used on your computer?
2. Where do you find current version information about your RHEL 8 installation?
3. Which command shows a list of kernel modules that currently are loaded?
4. Which command enables you to discover kernel module parameters?
5. How do you unload a kernel module?
6. What can you do if you get an error message while trying to unload a kernel module?
7. How do you find which kernel module parameters are supported?
8. Where do you specify kernel module parameters that should be used persistently?
9. Assuming that the `cdrom` module has a parameter **debug**, which must be set to 1 to enable debug mode, which line would you include in the file that will automatically load that module?
10. How do you install a new version of the kernel?

End-of-Chapter Lab

In the end-of-chapter lab, you install a new version of the kernel and work with kernel modules.

Lab 16.1

1. Find out whether a new version of the kernel is available. If so, install it and reboot your computer so that it is used.
2. Use the appropriate command to show recent events that have been logged by the kernel.
3. Locate the kernel module that is used by your network card. Find out whether it has options. Try loading one of these kernel module options manually; if that succeeds, take the required measures to load this option persistently.



The following topics are covered in this chapter:

- Managing Systemd Targets
- Working with GRUB 2

The following RHCSA exam objectives are covered in this chapter:

- Configure systems to boot into a specific target automatically
- Modify the system bootloader

In this chapter, you learn how the boot procedure on Red Hat Enterprise Linux is organized. In the first part of this chapter, you learn about Systemd targets and how you can use them to boot your Linux system into a specific state. The second part of this chapter discusses GRUB2 and how to apply changes to the GRUB 2 boot loader. Troubleshooting is not a topic in this chapter; it is covered in Chapter 18, “Essential Troubleshooting Skills.”

Managing and Understanding the Boot Procedure

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 17-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 17-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Working with Systemd	1–7
Working with GRUB 2	8–10

1. Which of the following is the most efficient way to define a system want?
 - a. Use the **systemctl enable** command.
 - b. Define the want in the unit file [Service] section.
 - c. Create a symbolic link in the /usr/lib/systemd/system directory.
 - d. Create a symbolic link in the unit wants directory in the /etc/systemd/system directory.
2. Which target is considered the normal target for servers to start in?
 - a. graphical.target
 - b. server.target
 - c. multi-user.target
 - d. default.target

3. Which of the following is *not* an example of a system target?
 - a. `rescue.target`
 - b. `restart.target`
 - c. `multi-user.target`
 - d. `graphical.target`
4. Where do you define which target a unit should be started in if it is enabled?
 - a. The target unit file
 - b. The wants directory
 - c. The `systemctl.conf` file
 - d. The `[Install]` section in the unit file
5. To allow targets to be isolated, you need a specific statement in the target unit file. Which of the following describes that statement?
 - a. **`AllowIsolate`**
 - b. **`Isolate`**
 - c. **`SetIsolate`**
 - d. **`Isolated`**
6. An administrator wants to change the current `multi-user.target` to the `rescue.target`. Which of the following should she do?
 - a. Use **`systemctl isolate rescue.target`**
 - b. Use **`systemctl start rescue.target`**
 - c. Restart the system, and from the GRUB boot prompt specify that `rescue.target` should be started
 - d. Use **`systemctl enable rescue.target --now`**
7. To which System V runlevel does `multi-user.target` correspond?
 - a. 2
 - b. 3
 - c. 4
 - d. 5

8. What is the name of the file where you should apply changes to the GRUB 2 configuration?
 - a. `/boot/grub/menu.lst`
 - b. `/boot/grub2/grub.cfg`
 - c. `/etc/sysconfig/grub`
 - d. `/etc/default/grub`
9. After applying changes to the GRUB 2 configuration, you need to write those changes. Which of the following commands will do that for you?
 - a. `grub2 -o /boot/grub/grub.cfg`
 - b. `grub2-mkconfig > /boot/grub2/grub.cfg`
 - c. `grub2 > /boot/grub2/grub.cfg`
 - d. `grub2-install > /boot/grub2/grub.cfg`
10. What is the name of the GRUB2 configuration file that is generated on a UEFI system?
 - a. `/boot/efi/redhat/grub.cfg`
 - b. `/boot/efi/EFI/redhat/grub.cfg`
 - c. `/boot/EFI/grub.cfg`
 - d. `/boot/EFI/efi/grub.cfg`

Foundation Topics

Managing Systemd Targets

Systemd is the service in Red Hat Enterprise Linux 8 that is responsible for starting all kinds of things. Systemd goes way beyond starting services; other items are started from Systemd as well. In Chapter 11, “Working with Systemd,” you learned about the Systemd fundamentals; this chapter looks at how Systemd targets are used to boot your system into a specific state.

Understanding Systemd Targets

A Systemd *target* is basically just a group of units that belong together. Some targets are just that and nothing else, whereas other targets can be used to define the state a system is booting in, because these targets have one specific property that regular targets don’t have: they can be isolated. Isolatable targets contain everything a system needs to boot or change its current state. Four targets can be used while booting:

- **emergency.target:** In this target only a minimal number of units are started, just enough to fix your system if something is seriously wrong. You’ll find that it is quite minimal, as some important units are not started.
- **rescue.target:** This target starts all units that are required to get a fully operational Linux system. It doesn’t start nonessential services though.
- **multi-user.target:** This target is often used as the default target a system starts in. It starts everything that is needed for full system functionality and is commonly used on servers.
- **graphical.target:** This target also is commonly used. It starts all units that are needed for full functionality, as well as a graphical interface.

Working with Targets

Working with targets may seem complicated, but it is not. It drills down to three common tasks:

- Adding units to be automatically started
- Setting a default target
- Running a nondefault target to enter troubleshooting mode

In Chapter 11 you learned how to use the **systemctl enable** and **systemctl disable** commands to add services to or remove services from targets. In this chapter you'll learn how to set a default target and how to run a nondefault target to enter troubleshooting mode. But first we'll take a closer look at the working of targets under the hood.

Understanding Target Units



Behind a target there is some configuration. This configuration consists of two parts:

- The target unit file
- The “wants” directory, which contains references to all unit files that need to be loaded when entering a specific target

Targets by themselves can have dependencies to other targets, which are defined in the target unit file. Example 17-1 shows the definition of the multi-user.target file, which defines the normal operational state of a RHEL server.

Example 17-1 The multi-user.target File

```
[root@localhost ~]# systemctl cat multi-user.target
# /usr/lib/systemd/system/multi-user.target
# SPDX-License-Identifier: LGPL-2.1+
#
# This file is part of systemd.
#
# systemd is free software; you can redistribute it and/or modify it
# under the terms of the GNU Lesser General Public License as
# published by
# the Free Software Foundation; either version 2.1 of the License,
# or
# (at your option) any later version.

```

```
[Unit]
Description=Multi-User System
Documentation=man:systemd.special(7)
Requires=basic.target
Conflicts=rescue.service rescue.target
After=basic.target rescue.service rescue.target
AllowIsolate=yes
```


You can see that by itself the target unit does not contain much. It just defines what it requires and which services and targets it cannot coexist with. It also defines load ordering, by using the **After** statement in the [Unit] section. The target file does not contain any information about the units that should be included; that is in the individual unit files and the wants (explained in the upcoming section “Understanding Wants”).

Systemd targets look a bit like runlevels used in previous versions of RHEL, but targets are more than that. A target is a group of units, and there are multiple different targets. Some targets, such as the `multi-user.target` and the `graphical.target`, define a specific state that the system needs to enter. Other targets just bundle a group of units together, such as the `nfs.target` and the `printer.target`. These targets are included from other targets, such as `multi-user.target` or `graphical.target`.

Understanding Wants

Understanding the concept of a want simply requires understanding the verb *want* in the English language, as in “I want a cookie.” Wants in Systemd define which units Systemd wants when starting a specific target. Wants are created when Systemd units are enabled using **systemctl enable**, and this happens by creating a symbolic link in the `/etc/systemd/system` directory. In this directory, you’ll find a subdirectory for every target, containing wants as symbolic links to specific services that are to be started.



Managing Systemd Targets

As an administrator, you need to make sure that the required services are started when your server boots. To do this, use the **systemctl enable** and **systemctl disable** commands. You do not have to think about the specific target a service has to be started in. Through the [Install] section in the service unit file, the services know for themselves in which targets they need to be started, and a want is created automatically in that target when the service is enabled. The following procedure walks you through the steps of enabling a service:

1. Type **systemctl status vsftpd**. If the service has not yet been enabled, the Loaded line will show that it currently is disabled:

```
[root@server202 ~]# systemctl status vsftpd
vsftpd.service - Vsftpd ftp daemon
    Loaded: loaded (/usr/lib/systemd/system/vsftpd.service; disabled)
    Active: inactive (dead)
```

2. Type **ls /etc/systemd/system/multi-user.target.wants**. You’ll see symbolic links that are taking care of starting the different services on your machine. You can also see that the `vsftpd.service` link does not exist.

3. Type **systemctl enable vsftpd**. The command shows you that it is creating a symbolic link for the file `/usr/lib/systemd/system/vsftpd.service` to the directory `/etc/systemd/system/multi-user.target.wants`. So basically, when you enable a Systemd unit file, in the background a symbolic link is created.

TIP On the RHCSA exam, you are likely to configure a couple of services. It is a good idea to read through the exam questions, identify the services that need to be enabled, and enable them all at once to make sure that they are started automatically when you restart. This prevents your being so focused on configuring the service that you completely forget to enable it as well.

Isolating Targets

As already discussed, on Systemd machines there are several targets. You also know that a target is a collection of units. Some of those targets have a special role because they can be isolated. These are also the targets that you can set as the targets to get into after system start.

By isolating a target, you start that target with all of its dependencies. Only targets that have the **isolate** option enabled can be isolated. We'll explore the **systemctl isolate** command later in this section. Before doing that, let's take a look at the default targets on your computer.

To get a list of all targets currently loaded, type **systemctl --type=target**. You'll see a list of all the targets currently active. If your server is running a graphical environment, this will include all the dependencies required to install the graphical.target also. However, this list shows only the active targets, not all the targets. Type **systemctl --type=target --all** for an overview of all targets that exist on your computer. You'll now see inactive targets also (see Example 17-2).

Example 17-2 Showing System Targets

```
root@localhost ~]# systemctl --type=target --all
```

UNIT	LOAD	ACTIVE	SUB	DESCRIPTION
basic.target	loaded	active	active	Basic System
bluetooth.target	loaded	active	active	Bluetooth
cryptsetup.target	loaded	active	active	Local Encrypted Volumes
dbus.target	not-found	inactive	dead	dbus.target
emergency.target	loaded	inactive	dead	Emergency Mode
getty-pre.target	loaded	active	active	Login Prompts (Pre)
getty.target	loaded	active	active	Login Prompts

graphical.target	loaded	active	active	Graphical Interface
initrd-fs.target	loaded	inactive	dead	Initrd File Systems
initrd-root-device.target	loaded	inactive	dead	Initrd Root Device
initrd-root-fs.target	loaded	inactive	dead	Initrd Root File System
initrd-switch-root.target	loaded	inactive	dead	Switch Root
initrd.target	loaded	inactive	dead	Initrd Default Target
local-fs-pre.target	loaded	active	active	Local File Systems (Pre)
local-fs.target	loaded	active	active	Local File Systems
multi-user.target	loaded	active	active	Multi-User System
network-online.target	loaded	active	active	Network is Online
network-pre.target	loaded	active	active	Network (Pre)
network.target	loaded	active	active	Network
nfs-client.target	loaded	active	active	NFS client services
nss-lookup.target	loaded	inactive	dead	Host and Network Name Lookups
nss-user-lookup.target	loaded	active	active	User and Group Name Lookups
paths.target	loaded	active	active	Paths
remote-fs-pre.target	loaded	active	active	Remote File Systems (Pre)
remote-fs.target	loaded	active	active	Remote File Systems
rescue.target	loaded	inactive	dead	Rescue Mode
rpc_pipefs.target	loaded	active	active	rpc_pipefs. target
rpcbind.target	loaded	active	active	RPC Port Mapper
shutdown.target	loaded	inactive	dead	Shutdown
slices.target	loaded	active	active	Slices
sockets.target	loaded	active	active	Sockets
sound.target	loaded	active	active	Sound Card
sshd-keygen.target	loaded	active	active	sshd-keygen. target
swap.target	loaded	active	active	Swap
sysinit.target	loaded	active	active	System Initialization

Of the targets on your system, a few have an important role because they can be started (isolated) to determine the state your server starts in. These are also the targets that can be set as the default targets. These targets also roughly correspond to runlevels used on earlier versions of RHEL. These are the following targets:

<code>poweroff.target</code>	runlevel 0
<code>rescue.target</code>	runlevel 1
<code>multi-user.target</code>	runlevel 3
<code>graphical.target</code>	runlevel 5
<code>reboot.target</code>	runlevel 6

If you look at the contents of each of these targets, you'll also see that they contain the `AllowIsolate=yes` line. That means that you can switch the current state of your computer to either one of these targets using the **`systemctl isolate`** command. Exercise 17-1 shows you how to do this.

Exercise 17-1 Isolating Targets

1. From a root shell, go to the directory `/usr/lib/systemd/system`. Type **`grep Isolate *.target`**. This shows a list of all targets that allow isolation.
2. Type **`systemctl isolate rescue.target`**. This switches your computer to `rescue.target`. You need to type the root password on the console of your server to log in.
3. Type **`systemctl isolate reboot.target`**. This restarts your computer.

Key
Topic

Setting the Default Target

Setting the default target is an easy procedure that can be accomplished from the command line. Type **`systemctl get-default`** to see the current default target and use **`systemctl set-default`** to set the desired default target.

To set the `graphical.target` as the default target, you need to make sure that the required packages are installed. If this is not the case, you can use the **`yum group list`** command to show a list of all RPM package groups. The “server with gui” and “GNOME Desktop” package groups both apply. Use **`yum group install "server with gui"`** to install all GUI packages on a server where they have not been installed yet.

Working with GRUB 2

The GRUB 2 boot loader is one of the first things that needs to be working well to boot a Linux server. As an administrator, you will sometimes need to apply modifications to the GRUB 2 boot loader configuration. This section explains how to do so. The RHEL 8 boot procedure is discussed in more detail in Chapter 18, where troubleshooting topics are covered as well.

Understanding GRUB 2

The GRUB 2 boot loader makes sure that you can boot Linux. GRUB 2 is installed in the boot sector of your server's hard drive and is configured to load a Linux kernel and the initramfs:

Key Topic

- The kernel is the heart of the operating system, allowing users to interact with the hardware that is installed in the server.
- The initramfs contains drivers that are needed to start your server. It contains a mini file system that is mounted during boot. In it are kernel modules that are needed during the rest of the boot process (for example, the LVM modules and SCSI modules for accessing disks that are not supported by default).

Normally, GRUB 2 works just fine and does not need much maintenance. In some cases, though, you might have to change its configuration. To apply changes to the GRUB 2 configuration, the starting point is the `/etc/default/grub` file, which has options that tell GRUB what to do and how to do it. Example 17-3 shows the contents of this file after an installation with default settings of RHEL 8.

Key Topic

Example 17-3 Contents of the `/etc/default/grub` File

```
[root@localhost ~]# cat /etc/default/grub
GRUB_TIMEOUT=5
GRUB_DISTRIBUTOR="$(sed 's, release .*$,,g' /etc/system-release)"
GRUB_DEFAULT=saved
GRUB_DISABLE_SUBMENU=true
GRUB_TERMINAL_OUTPUT="console"
GRUB_CMDLINE_LINUX="crashkernel=auto resume=/dev/mapper/rhel-swap
rd.lvm.lv=rhel/root rd.lvm.lv=rhel/swap rhgb quiet"
GRUB_DISABLE_RECOVERY="true"
GRUB_ENABLE_BLSCFG=true
```

As you can see, the `/etc/default/grub` file does not contain much information. The most important part that it configures is the `GRUB_CMDLINE_LINUX` option. This line contains boot arguments for the kernel on your server.

TIP For the RHCSA exam, make sure that you understand the contents of the `/etc/default/grub` file. That is the most important part of the GRUB 2 configuration anyway.

Apart from the configuration in `/etc/default/grub`, there are a few configuration files in `/etc/grub.d`. In these files, you'll find rather complicated shell code that tells GRUB what to load and how to load it. You typically do not have to modify these files. You also do not need to modify anything if you want the capability to select from different kernels while booting. GRUB 2 picks up new kernels automatically and adds them to the boot menu automatically, so nothing has to be added manually.

Understanding GRUB 2 Configuration Files

Based on the configuration files mentioned previously, the main configuration file is created. If your system is a BIOS system, the name of the file is `/boot/grub2/grub.cf`. On a UEFI system the file is written to `/boot/efi/EFI/redhat`. After making modifications to the GRUB 2 configuration, you'll need to regenerate the relevant configuration file, which is why you should know the name of the file that applies to your system architecture. Do *not* edit it, as this file is automatically generated.

Modifying Default GRUB 2 Boot Options

To apply modifications to the GRUB 2 boot loader, the file `/etc/default/grub` is your entry point. The most important line in this file is `GRUB_CMDLINE_LINUX`, which defines how the Linux kernel should be started. In this line, you can apply permanent fixes to the GRUB 2 configuration. Some likely candidates for removal are the options **rhgb** and **quiet**. These options tell the kernel to hide all output while booting. That is nice to hide confusing messages for end users, but if you are a server administrator, you probably just want to remove these options.

TIP On the exam, you want to know immediately if something does not work out well. To accomplish this, it is a good idea to remove the **rhgb** and **quiet** boot options. Without these you will not have to guess why your server takes a long time after a restart; you'll just be able to see.

Another interesting parameter is `GRUB_TIMEOUT`. This defines the amount of time your server waits for you to access the GRUB 2 boot menu before it continues booting automatically. If your server runs on physical hardware that takes a long time to get through the BIOS checks, it may be interesting to increase this time a bit.

While working with GRUB 2, you need to know a bit about kernel boot arguments. There are many of them, and most of them you'll never use, but it is good to know where you can find them. Type **man 7 bootparam** for a man page that contains an excellent description of all boot parameters that you may use while starting the kernel.

To write the modified configuration to the appropriate files, you use the **grub2-mkconfig** command and redirect its output to the appropriate configuration file. On a BIOS system, the command would be **grub2-mkconfig -o /boot/grub2/grub.cfg**, and on a UEFI system the command would be **grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg**.

In Exercise 17-2, you learn how to apply modifications to the GRUB 2 configuration and write them to the `/boot/grub2/grub.cfg` configuration file.

TIP You should know how to apply changes to the GRUB configuration, but you should also know that the default GRUB 2 configuration works fine as it is for almost all computers. So, you will probably never have to apply any changes at all!

Key Topic

Exercise 17-2 Applying Modifications to GRUB 2

1. Open the file `/etc/default/grub` with an editor and remove the **rhgb** and **quiet** options from the `GRUB_CMDLINE_LINUX` line.
2. From the same file, set the `GRUB_TIMEOUT` parameter to 10 seconds. Save changes to the file and close the editor.
3. From the command line, type **grub2-mkconfig > /boot/grub2/grub.cfg** to write the changes to GRUB 2. (Note that instead of using the redirector `>` to write changes to the `grub.cfg` file, you could use the **-o** option. Both methods have the same result.)
4. Reboot and verify that while booting you see boot messages scrolling by.

Summary

In this chapter you learned how Systemd and GRUB 2 are used to bring your server into the exact state you desire at the end of the boot procedure. You also learned how Systemd is organized, and how units can be configured for automatic start with the use of targets. Finally, you read how to apply changes to the default GRUB 2 boot loader. In the next chapter, you learn how to troubleshoot the boot procedure and fix some common problems.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 17-2 lists a reference of these key topics and the page number on which each is found.

Table 17-2 Key Topics for Chapter 17

Key Topic Element	Description	Page
Section	Understanding target units	389
Section	Managing Systemd targets	390
Exercise 17-1	Isolating targets	393
List	Explanation of the role of kernel and initramfs	394
Example 17-3	Contents of the /etc/default/grub file	394
Exercise 17-2	Applying modifications to GRUB 2	396



Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

unit, wants, target, Systemd, dependencies, initramfs, kernel, boot loader, GRUB

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. What is a unit?
2. Which command enables you to make sure that a target is no longer eligible for automatic start on system boot?

3. Which configuration file should you modify to apply common changes to GRUB 2?
4. Which command should you use to show all service units that are currently loaded?
5. How do you create a want for a service?
6. How do you switch the current operational target to the rescue.target?
7. Why can it happen that you get the message that a target cannot be isolated?
8. You want to shut down a Systemd service, but before doing that you want to know which other units have dependencies to this service. Which command would you use?
9. What is the name of the GRUB 2 configuration file where you apply changes to GRUB 2?
10. After applying changes to the GRUB 2 configuration, which command should you run?

End-of-Chapter Labs

You have now learned how to work with Systemd targets and the GRUB 2 boot loader. Before continuing, it is a good idea to work on some labs that help you ensure that you can apply the skills that you acquired in this chapter.

Lab 17.1

1. Set the default target to multi-user.target.
2. Reboot to verify this is working as expected.

Lab 17.2

1. Change your GRUB 2 boot configuration so that you will see boot messages upon startup.

This page intentionally left blank



The following topics are covered in this chapter:

- Understanding the RHEL 8 Boot Procedure
- Passing Kernel Boot Arguments
- Using a Rescue Disk
- Fixing Common Issues
- Recovering Access to a Virtual Machine

The following RHCSA exam objectives are covered in this chapter:

- Boot systems into different targets manually
- Interrupt the boot process in order to gain access to a system

In Chapter 17, “Managing and Understanding the Boot Procedure,” you learned how a RHEL 8 server boots and which role the boot loader GRUB 2 and Systemd play in that process. In this chapter, you learn what you can do when common problems occur while booting your server. This chapter teaches general approaches that help to fix some of the most common problems that may occur while booting. Make sure to master the topics discussed in this chapter well; they might save your (professional) life one day!

Essential Troubleshooting Skills

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 18-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 18-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Understanding the RHEL 8 Boot Procedure	1
Passing Kernel Boot Arguments	2–6
Using a Rescue Disk	7
Fixing Common Issues	8–9
Recovering Access to a Virtual Machine	10

1. Which of the following comes first in the Red Hat Enterprise Linux 8 boot procedure?
 - a. Systemd
 - b. Kernel
 - c. GRUB 2
 - d. Initramfs

2. You have just entered a kernel argument on the GRUB 2 boot prompt. Pressing which key(s) enables you to start with this boot argument?
 - a. ZZ
 - b. Ctrl-X
 - c. Esc
 - d. Enter
3. Your initramfs seems faulty and cannot initialize the LVM volumes on your disk. Which configuration file should you check for options that are used?
 - a. /etc/dracut.d/dracut.conf
 - b. /etc/dracut.conf
 - c. /etc/sysconfig/dracut
 - d. /etc/mkinitrd.conf
4. You do not have the root password and want to reset it. Which kernel argument offers the recommended way to reset it?
 - a. **init=/bin/bash**
 - b. **init=/bin/sh**
 - c. **systemd.unit=emergency.target**
 - d. **rd.break**
5. You want to see exactly what is happening on system boot. Which two boot options should you remove from the GRUB 2 boot prompt? (Choose two.)
 - a. **rhgb**
 - b. **logo**
 - c. **quiet**
 - d. **silent**
6. You want to enter the most minimal troubleshooting mode where as few services as possible are loaded. Which boot argument should you use?
 - a. **systemd.unit=break.target**
 - b. **systemd.unit=emergency.target**
 - c. **systemd.unit=rescue.target**
 - d. **1**

7. Which of the following situations can be resolved only by using a rescue disk?
 - a. The kernel stops loading.
 - b. The initramfs stops loading.
 - c. You never get to a GRUB 2 boot prompt.
 - d. You are prompted to enter the root password for maintenance mode.
8. You have entered a troubleshooting mode, and disk access is read-only. What should you do?
 - a. Restart the troubleshooting mode and pass the **rw** boot option to the kernel.
 - b. Use the **rd.break** boot argument to manually start into the initramfs mode.
 - c. Use **mount -o remount,rw /**.
 - d. Use **mount /**.
9. Your server shows a blinking cursor only while booting. No GRUB 2 menu is available. What is the first step in troubleshooting this issue?
 - a. From a rescue disk, try the **Boot from local disk** option.
 - b. Start a rescue environment and reinstall GRUB.
 - c. Start a rescue environment and re-create the initramfs.
 - d. Use the **rd.break** boot argument.
10. When recovering access to a virtual machine, you need to make the storage devices in the image file available. Which of the following commands would do that, assuming that the name of the image file is `/home/user/lab1.img`?
 - a. **partx -ax /home/user/lab1.img**
 - b. **kpartx -ax /home/user/lab1.img**
 - c. **kpartx -av /home/user/lab1.img**
 - d. **partx -av /home/user/lab1.img**.

Foundation Topics

Understanding the RHEL 8 Boot Procedure

To fix boot issues, it is essential to have a good understanding of the boot procedure. If an issue occurs during boot, you need to be able to judge in which phase of the boot procedure the issue occurs so that you can select the appropriate tool to fix the issue.

The following steps summarize how the boot procedure happens on Linux:

Key Topic

1. **Performing POST:** The machine is powered on. From the system firmware, which can be the modern Universal Extended Firmware Interface (UEFI) or the classical Basic Input Output System (BIOS), the Power-On Self-Test (POST) is executed, and the hardware that is required to start the system is initialized.
2. **Selecting the bootable device:** Either from the UEFI boot firmware or from the BIOS, a bootable device is located.
3. **Loading the boot loader:** From the bootable device, a boot loader is located. On RHEL, this is usually GRUB 2.
4. **Loading the kernel:** The boot loader may present a boot menu to the user or can be configured to automatically start a default operating system. To load Linux, the kernel is loaded together with the `initramfs`. The `initramfs` contains kernel modules for all hardware that is required to boot, as well as the initial scripts required to proceed to the next stage of booting. On RHEL 8, the `initramfs` contains a complete operational system (which may be used for troubleshooting purposes).
5. **Starting `/sbin/init`:** Once the kernel is loaded into memory, the first of all processes is loaded, but still from the `initramfs`. This is the `/sbin/init` process, which on RHEL is linked to `Systemd`. The `udev` daemon is loaded as well to take care of further hardware initialization. All this is still happening from the `initramfs` image.
6. **Processing `initrd.target`:** The `Systemd` process executes all units from the `initrd.target`, which prepares a minimal operating environment, where the root file system on disk is mounted on the `/sysroot` directory. At this point, enough is loaded to pass to the system installation that was written to the hard drive.
7. **Switching to the root file system:** The system switches to the root file system that is on disk and at this point can load the `Systemd` process from disk as well.

- 8. Running the default target:** Systemd looks for the default target to execute and runs all of its units. In this process, a login screen is presented, and the user can authenticate. Note that the login prompt can be prompted before all Systemd unit files have been loaded successfully. So, seeing a login prompt does not necessarily mean that your server is fully operational yet.

In each of the phases listed, issues may occur because of misconfiguration or other problems. Table 18-2 summarizes where a specific phase is configured and what you can do to troubleshoot if something goes wrong.

TIP Troubleshooting has always been a part of the RHCSA exam. If you encounter an issue, make sure that you can identify in which phase of the boot procedure it occurs and what you can do to fix it.

Table 18-2 Boot Phase Configuration and Troubleshooting Overview

**Key
Topic**

Boot Phase	Configuring It	Fixing It
POST	Hardware configuration (F2, Esc, F10, or another key).	Replace hardware.
Selecting the bootable device	BIOS/UEFI configuration or hardware boot menu.	Replace hardware or use rescue system.
Loading the boot loader	grub2-install and edits to <code>/etc/defaults/grub</code> .	GRUB boot prompt and edits to <code>/etc/defaults/grub</code> , followed by grub2-mkconfig .
Loading the kernel	Edits to the GRUB configuration and <code>/etc/dracut.conf</code> .	GRUB boot prompt and edits to <code>/etc/defaults/grub</code> , followed by <code>grub2-mkconfig</code> .
Starting <code>/sbin/init</code>	Compiled into <code>initramfs</code> .	init= kernel boot argument, rd.break kernel boot argument.
Processing <code>initrd.target</code>	Compiled into <code>initramfs</code> .	Not typically required.
Switch to the root file system	Edits to the <code>/etc/fstab</code> file.	Edits to the <code>/etc/fstab</code> file.
Running the default target	Using systemctl set-default to create the <code>/etc/systemd/system/default.target</code> symbolic link	Start the <code>rescue.target</code> as a kernel boot argument.

In the next section you learn how to apply the different troubleshooting techniques described in this table.

Passing Kernel Boot Arguments

If your server does not boot normally, the GRUB boot prompt offers a convenient way to stop the boot procedure and pass specific options to the kernel while booting. In this section, you learn how to access the boot prompt and how to pass specific boot arguments to the kernel while booting.

Accessing the Boot Prompt

When your server boots, you briefly see the GRUB 2 menu. Look fast because it will only last for a few seconds. From this boot menu you can type **e** to enter a mode where you can edit commands, or **c** to enter a full GRUB command prompt, as shown in Figure 18-1. To pass boot options to a starting kernel, use **e**.




FIGURE 18-1 Entering the GRUB Boot Prompt

After passing an **e** to the GRUB boot menu, you'll see the interface that is shown in Figure 18-2. From this interface, scroll down to locate the section that begins with `linux ($root)/vmlinuz` followed by a lot of arguments. This is the line that tells GRUB how to start a kernel, and by default it looks like this:

```
linux ($root)/vmlinuz-4.18.0-80.el8.x86_64 root=/dev/mapper/rhel-root
ro crash kernel=auto resume=/dev/mapper/rhel-swap rd.lvm.lv=rhel/
root rd.lvm.lv=rhel/swap rhgb quiet
```

```
load_video
set gfx_payload=keep
insmod gzio
linux ($root)/vmlinuz-4.18.0-80.el8.x86_64 root=/dev/mapper/rhel-root ro crash\
kernel=auto resume=/dev/mapper/rhel-swap rd.lvm.lv=rhel/root rd.lvm.lv=rhel/sw\
ap rhgb quiet
initrd ($root)/initramfs-4.18.0-80.el8.x86_64.img $tuned_initrd

Press Ctrl-x to start, Ctrl-c for a command prompt or Escape to
discard edits and return to the menu. Pressing Tab lists
possible completions.
```

FIGURE 18-2 Enter Boot Arguments on the Line That Starts with linux

To start, it is a good idea to remove the **rhgb** and **quiet** parts from this line; these arguments hide boot messages for you, and typically you do want to see what is happening while booting. In the next section you learn about some troubleshooting options that you can enter from the GRUB boot prompt.

After entering the boot options you want to use, press Ctrl-X to start the kernel with these options. Notice that these options are used one time only and are not persistent. To make them persistent you must modify the contents of the `/etc/default/grub` configuration file and use **grub2-mkconfig -o /boot/grub2/grub.cfg** to apply the modification. (Refer to Chapter 17 for more details about this procedure.)

Starting a Troubleshooting Target

If you encounter trouble when booting your server, you have several options that you can enter on the GRUB boot prompt:

- **rd.break:** This stops the boot procedure while still in the `initramfs` stage. This option is useful if you do not have the root password available. The complete procedure for recovering a missing root password follows later in this chapter.
- **init=/bin/sh or init=/bin/bash:** This specifies that a shell should be started immediately after loading the kernel and `initrd`. This is a useful option, but it's not the best option because in some cases you'll lose console access or miss other functionality.
- **systemd.unit=emergency.target:** This enters a mode that loads a bare minimum number of Systemd units. It requires a root password. To see that

**Key
Topic**

only a very limited number of unit files have been loaded, you can type the **systemctl list-units** command.

- **systemd.unit=rescue.target:** This starts some more Systemd units to bring you in a more complete operational mode. It does require a root password. To see that only a very limited number of unit files have been loaded, you can type the **systemctl list-units** command.

In Exercise 18-1, you learn how to enter the rescue.target mode. The other modes listed here are discussed in the following sections.

Exercise 18-1 Exploring Troubleshooting Targets

1. (Re)start your computer. When the GRUB menu shows, select the first line in the menu and press **e**.
2. Scroll down to the line that starts with `linux $(root)/vmlinuz`. At the end of this line, type **systemd.unit=rescue.target**. Also remove the options **rhgb quit** from this line.
3. Enter the root password when you are prompted for it.
4. Type **systemctl list-units**. This shows all unit files that are currently loaded. You can see that a basic system environment has been loaded.
5. Type **systemctl show-environment**. This shows current shell environment variables.
6. Type **systemctl reboot** to reboot your machine.
7. When the GRUB menu appears, press **e** again to enter the editor mode. At the end of the line that loads the kernel, type **systemd.unit=emergency.target**.
8. When prompted for it, enter the root password to log in.
9. After successful login, type **systemctl list-units**. Notice that the number of unit files loaded is reduced to a bare minimum.

Using a Rescue Disk

If you are lucky when you encounter trouble, you'll still be able to boot from hard disk. If you are a bit less lucky, you'll just see a blinking cursor on a system that does not boot at all. If that happens, you need a rescue disk. The default rescue image for Red Hat Enterprise Linux is on the installation disk. When booting from the installation disk, you'll see a Troubleshooting menu item. Select this item to get access to the options you need to repair your machine.

Restoring System Access Using a Rescue Disk

After selecting the Troubleshooting option, you are presented with the following options, as shown in Figure 18-3:

- **Install Red Hat Enterprise Linux 8 in Basic Graphics Mode:** This option reinstalls your machine. Do not use it unless you want to troubleshoot a situation where a normal installation does not work and you need a basic graphics mode. Normally, you should never need to use this option to troubleshoot a broken installation.
- **Rescue a Red Hat Enterprise Linux System:** This is the most flexible rescue system. In Exercise 18-2, you can explore it in detail. This should be the first option of choice when using a rescue disk.
- **Run a Memory Test:** Run this option if you encounter memory errors. It allows you to mark bad memory chips so that your machine can boot normally.
- **Boot from Local Drive:** If you cannot boot from GRUB on your hard disk, try this option first. It offers a boot loader that tries to install from your machine's hard drive, and as such is the least intrusive option available.




FIGURE 18-3 Starting from a Rescue Disk

After starting a rescue system, you usually need to enable full access to the on-disk installation. Typically, the rescue disk detects your installation and mounts it on the `/mnt/sysimage` directory. To fix access to the configuration files and their default locations as they should be available on disk, use the **chroot /mnt/sysimage** command to make the contents of this directory your actual working environment. If you do not use this **chroot** command, many utilities will not work, because if they write to a configuration file, it will be the version that exists on the read-only disk.

Using the **chroot** command ensures that all path references to configuration files are correct.

In Exercise 18-2, you learn how to use the Rescue a Red Hat Enterprise Linux System option to troubleshoot a system that does not boot anymore.

Exercise 18-2 Using the Rescue Option

1. Restart your server from the installation disk. Select the **Troubleshooting** menu option.
2. From the Troubleshooting menu, select **Rescue a Red Hat Enterprise Linux System**. This prompts you to press Enter to start the installation. Do not worry: This option does not overwrite your current configuration; it just loads a rescue system.
3. The rescue system now prompts you that it will try to find an installed Linux system and mount on /mnt/sysimage. Press 1 to accept the **Continue** option (see Figure 18-4).

```
Starting installer, one moment...
anaconda 29.19.0.40-1.el8 for Red Hat Enterprise Linux 8.0.0 started.
* installation log files are stored in /tmp during the installation
* shell is available on TTY2
* when reporting a bug add logs from /tmp as separate text/plain attachments
=====
Rescue

The rescue environment will now attempt to find your Linux installation and
mount it under the directory : /mnt/sysimage. You can then make any changes
required to your system. Choose '1' to proceed with this step.
You can choose to mount your file systems read-only instead of read-write by
choosing '2'.
If for some reason this process does not work choose '3' to skip directly to a
shell.

1) Continue
2) Read-only mount
3) Skip to shell
4) Quit (Reboot)

Please make a selection from the above: _

[anaconda11:main* 2:shell 3:log 4:storage-log 5:program-log Switch tab: Alt+Tab | Help: F1]
```

FIGURE 18-4 The rescue system looks for an installed system image and mounts it for you.

4. If a valid Red Hat installation was found, you are prompted that your system has been mounted under `/mnt/sysimage`. At this point, you can press Enter twice to access the rescue shell.
5. Your Linux installation at this point is accessible through the `/mnt/sysimage` directory. Type **`chroot /mnt/sysimage`**. At this point, you have access to your root file system and you can access all tools that you need to repair access to your system.
6. Type **`exit`** and **`reboot`** to restart your machine in a normal mode.

Reinstalling GRUB Using a Rescue Disk

One of the common reasons you might need to start a rescue disk is that the GRUB 2 boot loader breaks. If that happens, you might need to reinstall it. After you have restored access to your server using a rescue disk, reinstalling GRUB 2 is not hard to do and consists of two steps:

- Make sure that you have made the contents of the `/mnt/sysimage` directory available to your current working environment, using **`chroot`** as described before.
- Use the **`grub2-install`** command, followed by the name of the device on which you want to reinstall GRUB 2. So on a KVM virtual machine, the command to use is **`grub2-install /dev/vda`**, and on a physical server or a VMware or Virtual Box virtual machine, it is **`grub2-install /dev/sda`**.

Re-Creating the Initramfs Using a Rescue Disk

Occasionally, the initramfs image may get damaged as well. If this happens, you cannot boot your server into normal operational mode. To repair the initramfs image after booting into the rescue environment, you can use the **`dracut`** command. If used with no arguments, this command creates a new initramfs for the kernel currently loaded.

Alternatively, you can use the **`dracut`** command with several options to make an initramfs for specific kernel environments. There is also a configuration file with the name `/etc/dracut.conf` that you can use to include specific options while re-creating the initramfs. The **`dracut`** configuration is dispersed over different locations:

- `/usr/lib/dracut/dracut.conf.d/*.conf` contains the system default configuration files.
- `/etc/dracut.conf.d` contains custom dracut configuration files.
- `/etc/dracut.conf` is used as the master configuration file.

Example 18-1 shows an example of the default dracut.conf configuration file.

Example 18-1 /etc/dracut.conf Default Contents

```
[root@server1 ~]# cat /etc/dracut.conf
# PUT YOUR CONFIG HERE OR IN separate files named *.conf
# in /etc/dracut.conf.d
# SEE man dracut.conf(5)

# Sample dracut config file

#logfile=/var/log/dracut.log
#fileloglvl=6

# Exact list of dracut modules to use. Modules not listed here are
#   not going
# to be included. If you only want to add some optional modules use
# add_dracutmodules option instead.
#dracutmodules+="

# dracut modules to omit
#omit_dracutmodules+="

# dracut modules to add to the default
#add_dracutmodules+="

# additional kernel modules to the default
#add_drivers+="

# list of kernel filesystem modules to be included in the generic
#   initramfs
#filesystems+="

# build initrd only to boot current hardware
#hostonly="yes"
#

# install local /etc/mdadm.conf
#mdadmconf="no"

# install local /etc/lvm/lvm.conf
#lvmconf="no"
```

```
# A list of fsck tools to install. If it is not specified, module's
# hardcoded
# default is used, currently: "umount mount /sbin/fsck* xfs_db
# xfs_check
# xfs_repair e2fsck jfs_fsck reiserfsck btrfsck". The installation is
# opportunistic, so non-existing tools are just ignored.
#fscks=""

# inhibit installation of any fsck tools
#nofscks="yes"

# mount / and /usr read-only by default
#ro_mnt="no"

# set the directory for temporary files
# default: /var/tmp
#tmpdir=/tmp
```

TIP According to the RHCSA objectives, you should not have to work with a rescue disk on the exam. However, as a Linux administrator, you should expect the unexpected, which is why it is a good idea to ensure that you can handle common as well as less common troubleshooting scenarios.

Fixing Common Issues

In one small chapter such as this, it is not possible to consider all the possible problems one might encounter when working with Linux. There are some problems, though, that are more likely to occur than others. In this section you learn about some of the more common problems.

Reinstalling GRUB 2

Boot loader code does not disappear just like that, but on occasion it can happen that the GRUB 2 boot code gets damaged. In that case, you better know how to reinstall GRUB 2. The exact approach depends on whether your server is still in a bootable state. If it is, reinstalling GRUB 2 is fairly easy. Just type **grub2-install** followed by the name of the device to which you want to install it. The command has many different options to fine-tune what exactly will be installed, but you probably will not need them because, by default, the command installs everything you need to make your system bootable again.

Reinstalling GRUB 2 becomes a little bit more complicated if your machine is in a nonbootable state. If that happens, you first need to start a rescue system and restore access to your server from the rescue system. (See Exercise 18-2 for the exact procedure for how to do that.) After mounting your server's file systems on `/mnt/sysimage` and using **chroot /mnt/sysimage** to make the mounted system image your root image, reinstalling is as easy as described previously: Just run **grub2-install** to install GRUB 2 to the desired installation device. So if you are in a KVM virtual machine, run **grub2-install /dev/vda**, and if you are on a physical disk, run **grub2-install /dev/sda**.

Fixing the Initramfs

In rare cases, it might happen that the initramfs gets damaged. If you analyze the boot procedure carefully, you will learn that you have a problem with the initramfs because you'll never see the root file system getting mounted on the root directory, nor will you see any Systemd units getting started. If you suspect that you are having a problem with the initramfs, it is easy to re-create it. To re-create it using all default settings (which is fine in most cases), you can just run the **dracut --force** command. (Without **--force**, the command will refuse to overwrite your existing initramfs.)

When running the **dracut** command, you can use the `/etc/dracut.conf` configuration file to specify what exactly is written to the initramfs. In this configuration file, you can see options like **lvmconf="no"** that can be used to switch specific features on or off. Use these options to make sure that you have all the required functionality in your initramfs.

Recovering from File System Issues

If you make a misconfiguration to your file system mounts, the boot procedure may just end with the message "Give root password for maintenance." This message is, in particular, generated by the **fsck** command that is trying to verify the integrity of the file systems in `/etc/fstab` while booting. If **fsck** fails, manual intervention is required that may result in this message during boot. Make sure that you know what to do when this happens to you!

TIP Make sure to master this topic very well. File system–related topics have a heavy weight in the RHCSA objectives, and it is likely that you will need to create partitions and/or logical volumes and put them in `/etc/fstab` for automatic mounting. That also makes it likely that something will go wrong, and if that happens on the exam, you better make sure that you know how to fix it!

If a device is referred to that does not exist, or if there is an error in the UUID that is used to mount the device, for example, Systemd waits first to see whether the device comes back by itself. If that does not happen, it gives the message “Give root password for maintenance” (see Figure 18-5). If that happens, you should by all means first enter the root password. Then you can type **journalctl -xb** as suggested to see whether relevant messages providing information about what is wrong are written to the journal. If the problem is file system oriented, type **mount -o remount,rw /** to make sure the root file system is mounted read/write and analyze what is wrong in the `/etc/fstab` file and fix it.

```
[ TIME ] Timed out waiting for device dev-sdg.device.
[DEPEND] Dependency failed for /mnt.
[DEPEND] Dependency failed for Local File Systems.
[DEPEND] Dependency failed for Mark the need to relabel after reboot.
Starting Restore /run/initramfs on shutdown...
[ OK ] Reached target Sound Card.
[ OK ] Reached target Bluetooth.
[ OK ] Reached target Network (Pre).
[ OK ] Reached target Timers.
[ OK ] Reached target Network.
[ OK ] Reached target Network is Online.
Starting Notify NFS peers of a restart...
[ OK ] Started Logout off all iSCSI sessions on shutdown.
[ OK ] Reached target User and Group Name Lookups.
[ OK ] Reached target Paths.
[ OK ] Reached target Sockets.
Starting Tell Plymouth To Write Out Runtime Data...
Starting Import network configuration from initramfs...
[ OK ] Started Restore /run/initramfs on shutdown.
[ OK ] Started Notify NFS peers of a restart.
[ OK ] Started Tell Plymouth To Write Out Runtime Data.
[ OK ] Started Import network configuration from initramfs.
Starting Create Volatile Files and Directories...
[ OK ] Started Emergency Shell.
[ OK ] Reached target Emergency Mode.
[ OK ] Started Create Volatile Files and Directories.
Mounting RPC Pipe File System...
Starting Security Auditing Service...
Starting RPC Bind...
[ OK ] Started RPC Bind.
[ OK ] Mounted RPC Pipe File System.
[ OK ] Reached target rpc_pipefs.target.
[ OK ] Reached target NFS client services.
[ OK ] Reached target Remote File Systems (Pre).
[ OK ] Reached target Remote File Systems.
Starting Crash recovery kernel arming...
[ OK ] Started Security Auditing Service.
Starting Update UTMP about System Boot/Shutdown...
[ OK ] Started Update UTMP about System Boot/Shutdown.
Starting Update UTMP about System Runlevel Changes...
[ OK ] Started Update UTMP about System Runlevel Changes.
You are in emergency mode. After logging in, type "journalctl -xb" to view
system logs, "systemctl reboot" to reboot, "systemctl default" or "exit"
to boot into default mode.
Give root password for maintenance
(or press Control-D to continue): _
```

FIGURE 18-5 If You See This, You Normally Have an `/etc/fstab` Issue

Resetting the Root Password

A common scenario for a Linux administrator is that the root password has gone missing. If that happens, you need to reset it. The only way to do that is by booting into minimal mode, which allows you to log in without entering a password. To do so, follow these steps:

Key Topic

1. On system boot, press **e** when the GRUB 2 boot menu is shown.
2. Enter **rd.break** as a boot argument to the line that loads the kernel and press Ctrl-X to boot with this option.
3. You'll now be dropped at the end of the boot stage where `initramfs` is loaded, just before a mount of the root file system on the directory `/`.
4. Type **mount -o remount,rw /sysroot** to get read/write access to the system image.
5. At this point, make the contents of the `/sysimage` directory your new root directory by typing **chroot /sysroot**.
6. Now you can enter **passwd** and set the new password for the user root.
7. Because at this very early boot stage SELinux has not been activated yet, the context type on `/etc/shadow` will be messed up. If you reboot at this point, no one will be able to log in. So you must make sure that the context type is set correctly. To do this, at this point you should load the SELinux policy by using **load_policy -i**. You'll find more information about SELinux in Chapter 22, "Managing SELinux."
8. Now you can manually set the correct context type to `/etc/shadow`. To do this, type **chcon -t shadow_t /etc/shadow**.
9. Reboot. You can now log in with the changed password for user root.

NOTE In the preceding procedure you have read how to use the **load_policy -i** and **chcon** commands to correct the labels on the `/etc/shadow` file. An alternative (and easier) method is to create a file with the name `/.autorelabel`, which will force SELinux to restore labels that are set on the entire file system.

Recovering Access to a Virtual Machine

A special case of troubleshooting is a situation that involves a virtual machine. If you have a problem in the virtual machine, and you cannot access it using Virtual Machine Manager or **virsh console**, you also can't connect to it using SSH.

However, there is still an option. The technique described here is a bit advanced, but because it might someday help you reestablish access to a virtual machine that would otherwise have been lost, it is worth knowing about.

TIP The procedure that is described here is relatively complicated, and you should not have to apply it unless you mess up severely during the exam. So instead of memorizing this procedure, just be very careful to prevent issues like this from happening!

1. To start, open a root shell on the KVM host. From that host file, make sure that the virtual machine is stopped by using **virsh destroy *vmname***. (Use the name of the virtual machine as listed by the **virsh list** command.)
2. Find the disk image file. It normally is stored in the `/var/lib/libvirt/images` directory. If you cannot find it there, use **virsh dumpxml *vmname* | grep "source file="** to find the name of the source file:

```
[root@lab ~]# virsh dumpxml sander-vm1 | grep "source file="
<source file='/home/sander/lab1.img' />
<source file='/var/lib/libvirt/images/sander-vm1.img' />
```

3. Now that you know the name of the virtual machine image file, you can mount it into the host environment by using the **kpartx -a** command. This command analyzes the disk layout on the virtual machine and creates storage devices that allow you to mount devices in the virtual machine in the host file system:

```
[root@lab ~]# kpartx -av /home/sander/lab1.img
add map loop0p1 (253:5): 0 1024000 linear /dev/loop0 2048
add map loop0p2 (253:6): 0 7362560 linear /dev/loop0 1026048
```

4. The **kpartx** command has created devices that enable you to access the two partitions in the virtual machine. If this is a typical RHEL 8 setup, the first partition is the `/boot` partition, and the second partition is normally used for LVM logical volumes. You can mount the `/boot` partition easily by using **mount /dev/mapper/loop0p1 /mnt**.
5. To access the logical volumes that exist within the virtual machine's second partition, use the **pvscan** command:

```
[root@lab mapper]# pvscan /dev/mapper/loop0p2
PV /dev/sda3 VG centos lvm2 [48.83 GiB / 0 free]
PV /dev/sda4 VG centos lvm2 [50.00 GiB / 20.00 GiB free]
PV /dev/sda5 VG vglvm lvm2 [347.32 GiB / 329.32 GiB free]
PV /dev/mapper/loop0p2 VG centosvm lvm2 [3.51 GiB / 0 free]
Total: 4 [449.65 GiB] / in use: 4 [449.65 GiB] / in no VG: 0 [0 ]
```

6. You now have activated the LVM setup within the virtual machine. If you type the **lvs** command, you should see the logical volumes listed as well, but in an active state. To activate them, use the **vgchange** command on the volume group that was found in the virtual machine disk image file. So for this exam, you type **vgchange -a y centosvm** to activate all logical volumes. After doing so, you can mount them also to directly access all files in the virtual machine file system and fix all issues that prevented the virtual machine from booting normally.

WARNING This procedure works well unless within the virtual machine the same LVM volume group name is used as on the host machine. This will normally not be the case because the VG name by default is the same as the hostname (and you normally want hostnames to be unique). If the VG names are not unique, though, you risk mixing them up and working on the wrong logical volume.

7. After you have fixed all issues in the virtual machine, you first need to unmount all file systems currently mounted. Next, you need to remove the device files that have been created based on the contents of the virtual machine image file by using the **kpartx -dv /home/sander/lab1.img** command. (Make sure to use the filename of the image file that is actually used on your computer.)

Summary

In this chapter, you learned how to troubleshoot the Red Hat Enterprise Linux 8 boot procedure. You learned in general what happens when a server boots and at which specific points you can interfere to fix things that go wrong. You also learned what to do in some specific cases. Make sure that you know these procedures well; you are likely to encounter them on the exam.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 18-3 lists a reference of these key topics and the page number on which each is found.

Table 18-3 Key Topics for Chapter 18

Key Topic Element	Description	Page
List	Summary of phases processed while booting	404
Table 18-2	Boot phase configuration and troubleshooting overview	405
List	Summary of relevant GRUB 2 boot options for troubleshooting	407
List	Resetting the root password	416

**Key
Topic**

Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

target, GRUB, initramfs, **dracut**

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. Which key do you need to press to enter the GRUB boot menu editor mode?
2. During startup, the boot procedure is not completed and the server asks for the root password instead. What is likely to be the reason of this?

3. You want to enter troubleshooting mode, but you do not know the root password. Which argument would you pass to the kernel to enter a mode that provides access to most of the machine's functionality?
4. You start your server and nothing happens. You just see a blinking cursor and that's all. What is the first step to troubleshoot this issue?
5. You want to find out which units are available in a specific troubleshooting environment. Which command would you use?
6. You want to start troubleshooting a lost root password. Which argument would you pass to the GRUB 2 boot loader?
7. From the shell that you have started to troubleshoot a lost password for user root, you want to load the SELinux policy. Which command enables you to do that?
8. While troubleshooting the root password, what do you need to do to make sure the SELinux labels are set correctly?
9. You have applied changes to the GRUB 2 boot loader and want to save them. How do you do that?
10. You do know the root password on a machine where you want to enter the most minimal troubleshooting mode. Which GRUB 2 boot argument would you use?

End-of-Chapter Lab

Lab 18.1 shows you how to troubleshoot some common problems.

Lab 18.1

1. Restart your server and change the root password from the appropriate troubleshooting mode.
2. In `/etc/fstab`, change one of the device names so that on next reboot the file system on it cannot be mounted. Restart and fix the issue that you encounter.
3. Use a rescue disk to bring your server up in full troubleshooting mode from the rescue disk.
4. Re-create the `initramfs`.

This page intentionally left blank



The following topics are covered in this chapter:

- Understanding Shell Scripting Core Elements
- Using Variables and Input
- Using Conditional Loops

Shell scripting is a science all by itself. You do not learn about all the nuts and bolts related to this science in this chapter. Instead, you learn how to apply basic shell scripting elements, which allows you to write a simple shell script and analyze what is happening in a shell script.

An Introduction to Bash Shell Scripting

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 19-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 19-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Understanding Shell Scripting Core Elements	1–2
Using Variables and Input	3–5
Using Conditional Loops	6–10

1. Which line should every Bash shell script start with?
 - a. `/bin/bash`
 - b. `#!/bin/bash`
 - c. `!#/bin/bash`
 - d. `!/bin/bash`
2. What is the purpose of the `exit 0` command that can be used at the end of a script?
 - a. It informs the parent shell that the script could be executed without any problems.
 - b. It makes sure the script can be stopped properly.
 - c. It is required only if a `for` loop has been used to close the `for` loop structure.
 - d. It is used to terminate a conditional structure in the script.

3. How do you stop a script to allow a user to provide input?
 - a. **pause**
 - b. **break**
 - c. **read**
 - d. **stop**
4. Which line stores the value of the first argument that was provided when starting a script in the variable NAME?
 - a. **NAME = \$1**
 - b. **\$1 = NAME**
 - c. **NAME = \$@**
 - d. **NAME=\$1**
5. What is the best way to distinguish between different arguments that have been passed into a shell script?
 - a. **\$?**
 - b. **\$#**
 - c. **\$***
 - d. **\$@**
6. What is used to close an **if** loop?
 - a. **end**
 - b. **exit**
 - c. **stop**
 - d. **fi**
7. What is missing in the following script at the position of the dots?

```
if [ -f $1 ]
then
    echo "$1 is a file"
..... [ -d $1 ]
then
    echo "$1 is a directory"
else
    echo "I do not know what \"$1\" is"
fi
```

- a. else
- b. if
- c. elif
- d. or

8. What is missing in the following script at the position of the dots?

```
for (( counter=100; counter>1; counter-- )); .....  
    echo $counter  
done  
exit 0
```

- a. in
- b. do
- c. run
- d. start

9. Which command is used to send a message with the subject “error” to the user root if something didn’t work out in a script?

- a. mail error root
- b. mail -s error root
- c. mail -s error root .
- d. mail -s error root < .

10. In a **case** statement, it is a good idea to include a line that applies to all other situations. Which of the following would do that?

- a. *)
- b. *
- c. else
- d. or

Foundation Topics

Understanding Shell Scripting Core Elements

Basically, a shell script is a list of commands that is sequentially executed, with some optional scripting logic in it that allows code to be executed under specific conditions only. To understand complex shell scripts, it is a good idea to start with an example of a very basic script, shown in Example 19-1.

Example 19-1 Basic Script Example

```
#!/bin/bash
#
# This is a script that greets the world
# Usage: ./hello

clear
echo hello world

exit 0
```

This basic script contains a few elements that should be used in all scripts. To start, there is the *shebang*. This is the line `#!/bin/bash`. When a script is started from a parent shell environment, it opens a subshell. In this subshell, different commands are executed. These commands can be interpreted in any way, and to make it clear how they should be interpreted, the shebang is used. In this case, the shebang makes clear that the script is a Bash shell script. Other shells can be specified as well. For instance, if your script contains Perl code, the shebang should be `#!/usr/bin/perl`.

It is good practice to start a script with a shebang; if it is omitted, the script code will be executed by the shell that is used in the parent shell as well. Because your scripts may also be executed by, for instance, users of `ksh`, using a shebang to call `/bin/bash` as a subshell is important to avoid confusion.

Right after the shebang, there is a part that explains what the script is about. It is a good idea in every script to include a few comment lines. In a short script, it is often obvious what the script is doing. If the script is becoming longer, and as more people get involved in writing and maintaining the script, it will often become less clear what the writer of the script intended to do. To avoid that, make sure that you include comment lines, starting with a `#`. Include them not only in the beginning of the script but also at the start of every subsection of the script. Comments will surely be helpful if you read your script a few months later and don't remember exactly

what you were trying to do while creating it. You can also use comments within lines. No matter in which position the `#` is used, everything from the `#` until the end of the line is comment text.

Next is the body of the script. In Example 19-1, the body is just a simple script containing a few commands that are sequentially executed. The body may grow as the script develops.

At the end of the script I have included the statement **exit 0**. An **exit** statement tells the parent shell whether the script was successful. A 0 means that it was successful, and anything else means that the script has encountered a problem. The exit status of the last command in the script is the exit status of the script itself, unless the **exit** command is used at the end of the script. But it is good to know that you can work with **exit** to inform the parent shell how it all went. To request the exit status of the last command, from the parent shell, use the command **echo \$?**. This can be useful to determine whether and why something didn't work out.

After writing a script, make sure that it can be executed. The most common way to do this is by applying the execute permission to it. So, if the name of the script is **hello**, use **chmod +x hello** to make it executable. The script can also be executed as an argument of the **bash** command, for instance. Use **bash hello** to run the **hello** script. If started as an argument of the **bash** command, the script does not need to be executable.

You can basically store the script anywhere you like, but if you are going to store it in a location that is not included in the `$PATH`, you need to execute it with a `./` in front of the script name. So, just typing **hello** is not enough to run your script; type **./hello** to run it. Note that this is also required if you want to run the script from the current directory, because on Linux the current directory is not included in the `$PATH` variable. Or put it in a standard directory that is included in the `$PATH` variable, like `/usr/local/bin`. In Exercise 19-1 you apply these skills and write a simple shell script.

Exercise 19-1 Writing a Simple Shell Script

1. Use **vim** to create a script with the name **hello** in your home directory.
2. Give this script the contents that you see in Example 19-1 and close it.
3. Use **./hello** to try to execute it. You get a “permission denied” error message.
4. Type **chmod +x hello** and try to execute it again. You see that it now works.

Using Variables and Input

Linux Bash scripts are much more than just a list of commands that is sequentially executed. One of the nice things about scripts is that they can work with variables and input to make the script flexible. In this section, you learn how to work with variables and input.

Using Positional Parameters

When starting a script, arguments can be used. An *argument* is anything that you put behind the script command while starting it. Arguments can be used to make a script more flexible. Take, for instance, the command **useradd lisa**. In this example, the command is **useradd**, and the argument **lisa** is specifying what needs to be done. In this case, a user with the name lisa has to be created. In this example, **lisa** is the argument to the command **useradd**. In a script, the first argument is referred to as **\$1**, the second argument is referred to as **\$2**, and so on. The example script in Example 19-2 shows how an argument can be used. Go ahead and try it using any arguments you want to use.

Example 19-2 Example Script That Is Using Arguments

```
#!/bin/bash
# run this script with a few arguments
echo The first argument is $1
echo The second argument is $2
echo The third argument is $3
```

If you tried to run the sample code from Example 19-2, you might have noticed that its contents are not perfect. If you use three arguments while using the script, it will work perfectly. If you only use two arguments, the third echo prints with no value for \$3. If you use four arguments, the fourth value (which would be stored in \$4) is never used. So, if you want to use arguments, you are better off using a more flexible approach. Example 19-3 shows an example of a script that is using a more flexible approach.

Example 19-3 Using Arguments in a Flexible Way

```
#!/bin/bash
# run this script with a few arguments
echo you have entered $# arguments
for i in "$@" do
    echo $i
done
exit 0
```

In Example 19-3, two new items that relate to the arguments are introduced:

- `$#` is a counter that shows how many arguments were used when starting the script.
- `$@` refers to all arguments that were used when starting the script.

To evaluate the arguments that were used when starting this script, a conditional loop with **for** can be used. In conditional loops with **for**, commands are executed as long as the condition is true. In this script, the condition is **for i in "\$@"**, which means “for each argument.” Each time the script goes through the loop, a value from the `$@` variable is assigned to the `$i` variable. So, as long as there are arguments, the body of the script is executed. The body of a **for** loop always starts with **do** and is closed with **done**, and between these two, the commands are listed that need to be executed. So, the example script in Example 19-3 will use **echo** to show the value of each argument and stop when no more arguments are available. In Exercise 19-2, you can try this for yourself by writing a script that works with positional parameters.

Exercise 19-2 Working with Positional Parameters

1. Open an editor, create a script named `ex192a`, and copy the contents from Example 19-2 into this script.
2. Save the script and make it executable.
3. Run the command `./ex192a a b c`. You see that three lines are echoed.
4. Run the command `./ex192a a b c d e f`. You see that still three lines are echoed.
5. Open an editor to create the script `ex192` and copy the contents from Example 19-3 into this script.
6. Save the script and make it executable.
7. Run the command `./ex192 a b c`. You see that three lines are echoed.
8. Run the command `./ex192` without arguments. You see that it does not echo anything.

Working with Variables

Key Topic

A *variable* is a label that is used to refer to a specific location in memory that contains a specific value. Variables can be defined statically by using `NAME=value` or in a dynamic way. There are two solutions to define a variable dynamically:

Key Topic

- Use **read** in the script to ask the user who runs the script for input.
- Use command substitution to use the result of a command and assign that to a variable. For example, the **date +%d-%m-%y** command shows the current date in day-month-year format. To assign that to a variable in a script, you could use **TODAY=\$(date +%d-%m-%y)**. In command substitution, you just have to enclose in parentheses the command whose result you want to use, with a dollar sign preceding the opening parentheses. As an alternative to this notation, backquotes may be used. So the command **TODAY=`date +%d-%m-%y`** would do exactly the same.

In the previous section about positional parameters, you learned how to provide arguments when starting a script. In some cases, it can be more efficient to ask for information when you find out that something essential is missing. The script in Example 19-4 shows how to do this using **read**.

Example 19-4 Example of a Script That Uses the **read** Command

```
#!/bin/bash
if [ -z $1 ]; then
    echo enter a name
    read NAME
else
    NAME=$1
fi
echo you have entered the text $NAME
exit 0
```

In Example 19-4, an **if ... then ... else ... fi** statement is used to check whether the argument **\$1** exists. This is done by using the **test** command, which can be written in either of two ways: **test** or **[...]**. In Example 19-4, the line **if [-z \$1]** executes to see if the test **-z \$1** is true. The **-z test** checks whether **\$1** is nonexistent. Stated otherwise, the line **if [-z \$1]** checks whether **\$1** is empty; if so, it means that no argument was provided when starting this script. If this is the case, the commands

after the **then** statement are executed. Notice that when writing the **test** command with the square brackets, it is essential to include one space after the opening bracket and one space before the closing bracket; without these spaces the command will not work.

Notice that the **then** statement immediately follows the **test** command. This is possible because a semicolon is used (;). A semicolon is a command separator and can replace a new line in a script. In the **then** statement, two commands are executed: an **echo** command that displays a message onscreen, and a **read** command. The **read** command stops the script so that user input can be processed and stored in the variable **NAME**. So, the line **read NAME** puts all user input in the variable **NAME**, which will be used later in the script.

In the example script in Example 19-4, the next part is introduced by the **else** statement. The commands after the **else** statement are executed in all other cases, which in this case means “if an argument was provided.” If that is the case, the variable **NAME** is defined and the current value of **\$1** is assigned to it.

Notice how the variable is defined: directly after the name of the variable there is an **=** sign, which is followed by **\$1**. Notice that you should never use spaces when defining variables.

Then, the **if** loop is closed with a **fi** statement. Once the **if** loop has been completed, you know for sure that the variable **NAME** is defined and has a value. The last line of the script reads the value of the variable **NAME** and displays this value to **STDOUT** via the **echo** command. Notice that to request the current value of a variable, the variable name is referred to, preceded by a **\$** sign.

In Exercise 19-3, you can practice working with input.

Exercise 19-3 Working with Input

1. Open an editor and create a script with the name **ex193**. Enter the contents of Example 19-4 in this script.
2. Write the script to disk and use **chmod +x ex193** to make it executable.
3. Run the script using **./ex193** and no further arguments. You see that it prompts for input.
4. Run the script using **hello** as its argument. It will echo “you have entered the text hello” to the **STDOUT**.

Using Conditional Loops

As you have already seen, you can use conditional loops in a script. These conditional loops are only executed if a certain condition is true. In Bash the following conditional loops are often used:

Key Topic

- **if ... then ... else:** Used to execute codes if a specific condition is true
- **for:** Used to execute commands for a range of values
- **while:** Used to execute code as long as a specific condition is true
- **until:** Used to execute code until a specific condition is true
- **case:** Used to evaluate specific values, where beforehand a limited number of values is expected

Working with if ... then ... else

The **if ... then ... else** construction is common to evaluate specific conditions. You have already seen an example with it in Example 19-4. This conditional loop is often used together with the **test** command, which you saw in action earlier to check whether a file exists. This command enables you to do many other things as well, such as compare files, compare integers, and much more.

TIP Take a look at the man page of the **test** command.

The basic construction with **if** is **if ... then ... fi**. This evaluates one single condition, as in the following example:

```
if [ -z $1 ]
then
    echo no value provided
fi
```

In Example 19-4 you saw how two conditions can be evaluated by including **else** in the statement. Example 19-5 shows how multiple conditions can be evaluated by contracting **else** with **if** to become **elif**. This is useful if many different values need to be checked. Note in Example 19-5 that multiple **test** commands are used as well.

Example 19-5 Example with **if ... then ... else**

```
#!/bin/bash
# run this script with one argument
# the goal is to find out if the argument is a file or a directory
if [ -f $1 ]
then
    echo "$1 is a file"
elif [ -d $1 ]
then
    echo "$1 is a directory"
else
    echo "I do not know what \"$1\" is"
fi
exit 0
```

Also note in Example 19-5 the use of the backslash (\). This character informs the shell that it should not interpret the following character, which is known as *escaping* the character.

Using || and &&

Instead of writing full **if ... then** statements, you can use the logical operators **||** and **&&**. **||** is a logical OR and will execute the second part of the statement only if the first part is not true; **&&** is the logical AND and will execute the second part of the statement only if the first part is true. Consider these two one-liners:

```
[ -z $1 ] && echo no argument provided
ping -c 1 10.0.0.20 2>/dev/null || echo node is not available
```

In the first example, a test is performed (using the alternative **test** command syntax) to see whether **\$1** is empty. If that test is true (which basically means that the **test** command exits with the exit code 0), the second command is executed.

In the second example, a **ping** command is used to check the availability of a host. The logical OR is used in this example to echo the text “node is not available” in case the **ping** command was not successful. You’ll often find that instead of fully written **if ... then** statements, the **&&** and **||** constructions are used. In Exercise 19-4 you can practice some **if ... then ... else** skills, using either **if ... then ... else** or **&&** and **||**.

Exercise 19-4 Using if ... then ... else

In this exercise, you work on a script that checks the availability of the Apache web server.

1. Start an editor and create a script with the name `filechk`.
2. Copy the contents from Example 19-5 to this script.
3. Run a couple of tests with it, such as `./filechk /etc/hosts`, `./filechk /usr`, and `./filechk non-existing-file`.

Applying for

The **for** conditional provides an excellent solution for processing ranges of data. In Example 19-6, you can see the first example with **for**, where a range is defined and processed as long as there are unprocessed values in that range.

Example 19-6 Example with for

```
#!/bin/bash
#
for (( COUNTER=100; COUNTER>1; COUNTER-- )); do
    echo $COUNTER
done
exit 0
```

A **for** conditional statement always starts with **for**, which is followed by the condition that needs to be checked. Then comes **do**, which is followed by the commands that need to be executed if the condition is true, and the conditional statement is closed with **done**.

In the example in Example 19-6, you can see that the condition is a range of numbers assigned to the variable `COUNTER`. The variable first is initialized with a value of 100, and as long as the value is higher than 1, in each iteration 1 is subtracted. As long as the condition is true, the value of the `$COUNTER` variable is displayed, using the **echo** commands.

In Example 19-7, you can see one of my favorite one-liners with **for**. The range is defined this time as a series of numbers, starting with 100 and moving up to 104.

Example 19-7 Example One-Liner with for

```
for i in {100..104}; do ping -c 1 192.168.4.$i >/dev/null && echo 192.168.4.$i is up; done
```

Notice how the range is defined: You specify the first number, followed by two dots and closed with the last number in the range. With **for i in**, each of these numbers is assigned to the variable **i**. For each of these numbers, a **ping** command is executed, where the option **-c 1** makes sure that one ping request only is sent.

In this **ping** command, it is not the result that counts, which is why the result is redirected to the **/dev/null** device. Based on the exit status of the **ping** command, the part behind the **&&** is executed. So, if the host could be reached, a line is echoed indicating that it is up.

Understanding while and until

Whereas the **for** statement that you have just read about is useful to work through ranges of items, the **while** statement is useful if you want to monitor something like the availability of a process. The counterpart of **while** is **until**, which keeps the iteration open until a specific condition is true. In Example 19-8 you can see how **while** is used to monitor process activity.

Example 19-8 Monitoring Processes with **while**

```
#!/bin/bash
#
# usage: monitor <processname>
while ps aux | grep $1 | grep -v grep > /dev/tty11
do
    sleep 5
done

clear
echo your process has stopped
logger $1 is no longer present
mail -s "process $1 has stopped" root < .
```

The script in Example 19-8 consists of two parts. First, there is the **while** loop. Second, there is everything that needs to be executed when the **while** loop no longer evaluates to true. The core of the **while** loop is the **ps** command, which is grepped for the occurrence of **\$1**. Notice the use of **grep -v grep**, which excludes lines containing the **grep** command from the result. Keep in mind that the **ps** command will include all running commands, including the **grep** command that the output of the **ps** command is piped to. This can result in a false positive match. The results of the **ps aux** command are redirected to **/dev/tty11**. That makes it possible to read the results later from **tty11** if that is needed, but they do not show by default.

After the **while** statements, follow the commands that need to be executed if the statement evaluates to true. In this case, the command is **sleep 5**, which will basically pause the script for 5 seconds. As long as the **while** command evaluates to true, it keeps on running. If it does no longer (which in this case means that the process is no longer available), it stops and the commands that follow the **while** loop can be executed.

In the line **mail -s “process \$1 has stopped” root < .**, a message is sent to the user root, using the internal mail handler that runs on RHEL 8 by default. The **mail** command takes as its first argument the subject, specified using the **-s** option. Notice the **< .** at the end of the command.

Normally, when using the **mail** command in an interactive mode, it will open an editor in which the message body can be written. This editor is closed by providing a line that has only a dot. In this command, the dot is provided through redirection of the STDIN. This allows the message to be processed without any further requirement for user activity.

The counterpart of **while** is **until**, which opens an iteration that lasts until the condition is true. In Example 19-9, **until** is used to filter the output of the **users** command for the occurrence of \$1, which would be a username. Until this command is true, the iteration continues. When the username has been found in the output of **users**, the iteration closes and the commands after the **until** loop are executed.

Example 19-9 Monitoring User Login with **until**

```
#!/bin/bash
#
until users | grep $1 > /dev/null
do
    echo $1 is not logged in yet
    sleep 5
done
echo $1 has just logged in
mail -s "$1 has just logged in" root < .
```

Understanding case

The last of the important iteration loops is **case**. The **case** statement is used to evaluate a number of expected values. The **case** statement in particular is important in Linux startup scripts that on previous versions of RHEL were used to start services. In a **case** statement, you define every specific argument that you expect, which is followed by the command that needs to be executed if that argument was used.

In Example 19-10, you can see the blueprint of the **case** statement that was used in the service scripts in earlier versions of RHEL to start almost any service. This statement works on \$1, which is the name of a startup script. Following the name of the script, the user can type **start**, **stop**, **restart**, and so on.

Example 19-10 Evaluating Specific Cases with **case**

```
case "$1" in
    start)
        start;;
    stop)
        rm -f $lockfile
        stop;;
    restart)
        restart;;
    reload)
        reload;;
    status)
        status
        ;;
    *)
        echo "Usage: $0 (start|stop|restart|reload|status)"
        ;;
esac
```

The **case** statement has a few particularities. To start, the generic syntax is **case** *item-to-evaluate* **in**. Then follows a list of all possible values that need to be evaluated. Each item is closed with a **)**. Then follows a list of commands that need to be executed if the specific argument was used. The list of commands is closed with a double semicolon. This **;;** can be used directly after the last command, and it can be used on a separate line. Also notice that the ***)** refers to all other options not previously specified. It is a “catchall” statement. The **case** iteration loop is closed by an **esac** statement.

Notice that the evaluations in **case** are performed in order. When the first match is made, the **case** statement will not evaluate anything else. Within the evaluation, wildcard-like patterns can be used. This shows in the ***)** evaluation, which matches everything. But you could as well use evaluations like **start|Start|START)** to match the use of a different case.

Bash Shell Script Debugging

When a script does not do what you expect it to do, debugging the script is useful. If a script does not do what you expect it to do, try starting it as an argument to the **bash -x** command. This will show you line by line what the script is trying to do, and also will show you specific errors if it does not work. Example 19-11 shows an example of using **bash -x** where it becomes immediately clear that the **grep** command does not know what it is expected to do, which is because it misses an argument to work on.

Example 19-11 Using **bash -x** to Debug Scripts

```
[root@server1 ~]# bash -x 319.sh
+ grep
Usage: grep [OPTION]... PATTERN [FILE]...
Try 'grep --help' for more information.
+ users
+ echo is not logged in yet
is not logged in yet
+ sleep 5
```

Summary

In this chapter you learned how to write shell scripts. You’ve worked through a few examples and are now familiar with some of the basic elements that are required to create a successful script.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 19-2 lists a reference of these key topics and the page number on which each is found.

Table 19-2 Key Topics for Chapter 19

Key Topic Element	Description	Page
Paragraph	Definition of variable	430
List	Dynamically defining variables	430
List	Conditional loops overview	432



Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

shebang, parent shell, subshell, variable, iteration, conditional loop, OR, AND

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. What is the effect if a script does *not* start with a shebang?
2. How can you check if a variable VAR has no value?
3. What would you use in a script to count the number of arguments that have been used?
4. What would you use to refer to all arguments that have been used when starting the script?
5. How do you process user input in a script?
6. What is the simplest way to test whether a file exists and execute the command **echo “file does not exist”** if it does not?
7. Which test would you perform to find out if an item is a file or a directory?
8. Which construction would you use to evaluate a range of items?
9. How do you close an **elif** statement in a script?
10. In a **case** statement, you evaluate a range of items. For each of these items you execute one or more commands. What do you need to use after the last command to close the specific item?

End-of-Chapter Lab

In this end-of-chapter lab, you apply your scripting skills to write two simple scripts.

Lab 19.1

1. Write a script that works with arguments. If the argument **one** is used, the script should create a file `/tmp/one`. If the argument **two** is used, the script should send a message containing the subject “two” to the root user.
2. Write a countdown script. The script should use one argument (and not more than one). This argument specifies the number of minutes to count down. It should start with that number of minutes and count down second by second, writing the text “there are nn seconds remaining” at every iteration. Use **sleep** to define the seconds. When there is no more time left, the script should echo “time is over” and quit.

This page intentionally left blank



The following topics are covered in this chapter:

- Hardening the SSH Server
- Using Other Useful sshd Options
- Configuring Key-Based Authentication with Passphrases

The following RHCSA exam objective is covered in this chapter:

- Configure key-based authentication for SSH

Secure Shell (SSH) is among the most important utilities that system administrators use. In Chapter 5, “Connecting to Red Hat Enterprise Linux 8,” you learned how to use SSH to connect to a server using a password or key-based authentication. In this chapter, you learn about some of the more advanced configuration settings.

Configuring SSH

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 20-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 20-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Hardening the SSH Server	1–5
Other Useful sshd Options	6–8, 10
Configuring Key-Based Authentication with Passphrases	9

1. Which of the following is *not* a common approach to prevent brute-force attacks against SSH servers?
 - a. Disable X11 forwarding
 - b. Have SSH listening on a nondefault port
 - c. Disable password login
 - d. Allow specific users only to log in
2. Which of the following successfully limits SSH server access to users bob and lisa only?
 - a. `LimitUsers bob,lisa`
 - b. `AllowedUsers bob lisa`
 - c. `AllowUsers bob lisa`
 - d. `AllowedUsers bob,lisa`

3. Which of the following commands must be used to provide nondefault port 2022 with the correct SELinux label?
 - a. `semanage ports -m -t ssh_port_t -p 2022`
 - b. `semanage port -m -t ssh_port_t -p tcp 2022`
 - c. `semanage ports -a -t sshd_port_t -p tcp 2022`
 - d. `semanage port -a -t ssh_port_t -p tcp 2022`
4. Which of the following descriptions is correct for the MaxAuthTries option?
 - a. After reaching the number of attempts specified here, the account will be locked.
 - b. This option specifies the maximum number of login attempts. After reaching half the number specified here, additional failures are logged.
 - c. After reaching the number of attempts specified here, the IP address where the login attempts come from is blocked.
 - d. The number specified here indicates the maximum number of login attempts per minute.
5. Which log file do you analyze to get information about failed SSH login attempts?
 - a. `/var/log/auth`
 - b. `/var/log/authentication`
 - c. `/var/log/messages`
 - d. `/var/log/secure`
6. SSH login in your test environment takes a long time. Which of the following options could be most likely responsible for the connection time problems?
 - a. UseLogin
 - b. GSSAPIAuthentication
 - c. UseDNS
 - d. TCPKeepAlive
7. Which of the following options is *not* used to keep SSH connections alive?
 - a. TCPKeepAlive
 - b. ClientAliveInterval
 - c. ClientAliveCountMax
 - d. UseDNS

8. Which file on an SSH client computer needs to be added to set the `ServerKeepAliveInterval` for an individual user?
 - a. `~/.ssh/ssh_config`
 - b. `~/.ssh/config`
 - c. `/etc/ssh/config`
 - d. `/etc/ssh/ssh_config`
9. Assuming that a passphrase-protected public/private key pair has already been created, how do you configure your session so that you have to enter the passphrase once only?
 - a. Copy the passphrase to the `~/.ssh/passphrase` file.
 - b. Run `ssh-add /bin/bash` followed by `ssh-agent`.
 - c. Run `ssh-agent /bin/bash` followed by `ssh-add`.
 - d. This is not possible; you must enter the passphrase each time a connection is created.
10. The `MaxSessions` option can be used to tune the maximum number of sessions that can be open at the same time. Which value does it have by default?
 - a. 10
 - b. 25
 - c. 100
 - d. 1000

Foundation Topics

Hardening the SSH Server

SSH is an important and convenient solution that helps you establish remote connections to servers. It is also a dangerous solution. If your SSH server is visible directly from the Internet, you can be sure that sooner or later an intruder will try to connect to your server, intending to do harm.

Dictionary attacks are common against an SSH server. The attacker uses the fact that SSH servers usually offer their services on port 22 and that every Linux server has a root account. Based on that information, it is easy for an attacker to try to log in as root just by guessing the password. If the password uses limited complexity, and no additional security measures have been taken, sooner or later the intruder will be able to connect. Fortunately, you can take some measures to protect SSH servers against these kinds of attacks:

- Disable root login
- Disable password login
- Configure a nondefault port for SSH to listen on
- Allow specific users only to log in on SSH

In the following subsections, you learn what is involved in changing these options.

Limiting Root Access

The fact that SSH servers by default have root login enabled is the biggest security problem. Disabling root login is easy; you just have to modify the `PermitRootLogin` parameter in `/etc/ssh/sshd_config` and reload or restart the service. After restarting, verify that you really cannot log in as root anymore.

TIP Some services automatically pick up changes in their configuration files. Most services pick up changes after the **`systemctl reload servicename`** command, whereas other services pick up changes only after a **`systemctl restart servicename`** command. To avoid wasting time on the exam, you should use **`systemctl restart servicename`** in all cases. At least you'll be sure that the service will pick up its new configuration.

After disabling root login, you must specify the username you want to use for login by using **`ssh user@servername`** or **`ssh -l user servername`**. If you do not specify the username, it takes the name of the current user on the client who is trying to open an SSH session.

Configuring Alternative Ports

Many security problems on Linux servers start with a port scan issued by the attacker. Scanning all of the 65,535 ports that can potentially be listening takes a lot of time, but most port scans focus on known ports only, and SSH port 22 is always among these ports. Do not underestimate the risk of port scans. On several occasions, I found that an SSH port listening at port 22 was discovered within an hour after installation of the server.

To protect against port scans, you can configure your SSH server to listen on another port. By default, `sshd_config` contains the line `Port 22` that tells SSH to listen on privileged port 22. To have SSH listen on another port, you must change port 22 into something else. Different ports can be used. You can choose to use a completely random port like 2022, but it can also be handy to configure SSH to listen on port 443.

Port 443 by default is assigned to web servers using Transport Layer Security (TLS) to offer encryption. If the users who want to access the SSH server are normally behind a proxy that allows traffic to ports 80 and 443 only, it may make sense to configure SSH to listen on port 443. You should realize, though, that by doing so port 443 cannot be used by your web server anymore; a port can be assigned to one service at a time only! So, do this only on a machine where you are not planning to run a TLS-enabled web server!

TIP To avoid being locked out of your server after making a change to the SSH listening port while being connected remotely, it is a good idea to open two sessions to your SSH server. Use one session to apply changes and test, and use the other session to keep your current connection open. Active sessions will not be disconnected after restarting the SSH server (unless you fail to restart the SSH server successfully).

Modifying SELinux to Allow for Port Changes

After changing the SSH port, you also need to configure SELinux to allow for this change. Network ports are labeled with SELinux security labels to prevent services from accessing ports where they should not go. To allow a service to connect to a nondefault port, you need to use **semanage port** to change the label on the target port. Before doing so, it is a good idea to check whether the port already has a label. You can do this by using the **semanage port -l** command.

If the port does not have a security label set yet, use **-a** to add a label to the port. If a security label has been set already, use **-m** to modify the current security label. Use,

for instance, the command **semanage port -a -t ssh_port_t -p tcp 2022** to label port 2022 for access by sshd. If you want to relabel a port that already was in use by another service, you have to use **semanage port -m** to modify the current port assignment. This is needed if, for instance, you want SSH to be able to bind to port 443.

Limiting User Access

Many options for sshd can be found by just browsing through the `sshd_config` file. One of the most interesting options to use is `AllowUsers`. This option takes a space-separated list of all users that will be allowed login through SSH. Notice that this is a powerful option, limiting login to only these users. If the user `root` still needs to be able to directly log in, you have to include `root` as well in the list of allowed users.

When using this parameter, it makes sense thinking about which username you want to allow or deny access. In a scripted brute-force attack, intruders normally also try common usernames such as `admin`, `Administrator`, and `jsmith`. It is easy to add a layer of security by selecting an uncommon username. Notice the following about the `AllowUsers` parameter:

- The `AllowUsers` option by default does not appear anywhere in the default `/etc/ssh/sshd_config` file.
- The `AllowUsers` option is a better option than `PermitRootLogin` because it is more restrictive than just denying `root` to log in.
- If the `AllowUsers` option does not specify `root`, you can still become `root` by using `su` - after making a connection as a normal user.

A parameter that looks promising, but is misleading, is `MaxAuthTries`. You might think that this option locks access to the SSH login prompt after a maximum number of failed login attempts. Such functionality proves useful when connecting to a local server (of which configuration can easily be changed if so required), but on an SSH server with Internet access, it is a rather dangerous option, making it easy to perform a denial-of-service attack on the server. An intruder would only have to run a script that tries to log in as a specific user to block access for that user for an amount of time. That is why `MaxAuthTries` does not do what you might think it would do. It just starts logging failed login attempts after half the number of successful login attempts specified here.

Still, the `MaxAuthTries` option is useful. For analyzing security events related to your SSH server, it is not that interesting to know when a user by accident has typed a wrong password one or two times. It becomes interesting only after multiple failed attempts. The higher the number of attempts, the more likely it is that an intruder is trying to get in. SSH writes log information about failed login attempts to the

AUTHPRIV syslog facility. By default, this facility is configured to write information about login failures to `/var/log/secure`.

In Exercise 20-1, you apply the common SSH options that have been discussed so far.

Exercise 20-1 Configuring SSH Security Options

In this exercise, the `sshd` process should be configured on `server1`. Use `server2` to test access to `server1`.

1. Open a root shell on `server1`, and from there, open the `sshd` configuration file `/etc/ssh/sshd_config` in an editor.
2. Find the `Port` line, and below that line add the line **Port 2022**. This tells the `sshd` process that it should bind to two different ports, which ensures that you can still open SSH sessions even if you have made an error.
3. Add the line **AllowUsers user** to the SSH configuration file as well.
4. Save changes to the configuration file and restart `sshd`, using **systemctl restart sshd**.
5. Type **systemctl status -l sshd**. You'll see a "permission denied" error for SSH trying to connect to port 2022.
6. Type **semanage port -a -t ssh_port_t -p tcp 2022** to apply the correct SELinux label to port 2022.
7. Open the firewall for port 2022 also, using **firewall-cmd --add-port=2022/tcp**, followed by **firewall-cmd --add-port=2022/tcp --permanent**.
8. Type **systemctl status -l sshd** again. You'll see that the `sshd` process is now listening on two ports.
9. Try to log in to your SSH server from your other server, using **ssh -p 2022 user@server1**. After the user shell has opened, type **su -** to get root access.

Using Other Useful `sshd` Options

Apart from the security-related options, there are some useful miscellaneous options that you can use to streamline SSH performance. In the next two subsections, you read about some of the most significant of these options.

Session Options

To start with, there is the `GSSAPIAuthentication` option, which on RHEL 8 is set to **yes** by default (which contradicts what the man page states about it). This option is useful in an environment where Kerberos authentication is used. If you do not have Kerberos in your environment, you might as well switch it off, because having this feature on slows down the authentication procedure.

The next interesting option is `UseDNS`. This option is on by default and instructs the SSH server to look up the remote hostname and check with DNS that the resolved hostname for the remote host maps back to the same IP address. Although this option has some security benefits, it also involves a significant performance penalty. If client connections are slow, make sure to set it to `no`, to switch off client hostname verification completely.

The third session-related option is `MaxSessions`. This specifies the maximum number of sessions that can be opened from one IP address simultaneously. If you are expecting multiple users to use the same IP address to log in to your SSH server, you might need to increase this option beyond its default value of 10.

Connection Keepalive Options

TCP connections in general are a relatively scarce resource, which is why connections that are not used for some time normally time out. You can use a few options to keep inactive connections alive for a longer period of time.

The `TCPKeepAlive` option is used to monitor whether the client is still available. Using this option (which is on by default) ensures that the connection is released for any machine that is inactive for a certain period of time. If used by itself, however, it might lead to a situation where unused connections are released as well, which is why it makes sense to use the `ClientAliveInterval` option. This option sets an interval, in seconds, after which the server sends a packet to the client if no activity has been detected. The `ClientAliveCountMax` parameter specifies how many of these packets should be sent. If the `ClientAliveInterval` is set to 30, and the `ClientAliveCountMax` is set to 10, for instance, inactive connections are kept alive for about 5 minutes. It is a good idea to set this to match the amount of time you want to keep inactive connections open.

The `ClientAliveInterval` and `ClientAliveCountMax` options can be specified on a server only. There is a client-side equivalent to these options also. If you cannot change the configuration of the SSH server, use the `ServerAliveInterval` and `ServerAliveCountMax` options to initiate connection keepalive traffic from the client machine. These options are set in the `/etc/ssh/ssh_config` file if they need to be applied for all users on that machine, or in `~/.ssh/config` if applied for individual users.

Table 20-2 provides an overview of the most useful SSH options.

Table 20-2 Most Useful sshd Configuration Options



Option	Use
Port	Defines the TCP listening port.
PermitRootLogin	Indicates whether to allow or disallow root login.
MaxAuthTries	Specifies the maximum number of authentication tries. After reaching half of this number, failures are logged to syslog.
MaxSessions	Indicates the maximum number of sessions that can be open from one IP address.
AllowUsers	Specifies a space-separated list of users who are allowed to connect to the server.
PasswordAuthentication	Specifies whether to allow password authentication. This option is on by default.
GSSAPIAuthentication	Indicates whether authentication through the GSSAPI needs to be enabled. Used for Kerberos-based authentication.
TCPKeepAlive	Specifies whether or not to clean up inactive TCP connections.
ClientAliveInterval	Specifies the interval, in seconds, that packets are sent to the client to figure out if the client is still alive.
ClientAliveCountMax	Specifies the number of client alive packets that need to be sent.
UseDNS	If on, uses DNS name lookup to match incoming IP addresses to names.
ServerAliveInterval	Specifies the interval, in seconds, that a client sends a packet to a server to keep connections alive.
ServerAliveCountMax	Specifies the maximum number of packets a client sends to a server to keep connections alive.

Configuring Key-Based Authentication with Passphrases

By default, password authentication is allowed on RHEL 8 SSH servers. If a public/private key pair is used, as explained in Chapter 5, this key pair is used first. If you want to allow public/private key-based authentication only and disable password-based authentication completely, set the PasswordAuthentication option to no.

When you use public/private keys, a passphrase can be used. Using a passphrase makes the key pair stronger. Not only does an intruder have to get access to the private key, but when he does, he must also know the passphrase to use the key. This is why for establishing client/server connections with public/private keys, it is recommended to use passphrases. Without further configuration, the use of passphrases would mean that users have to enter the passphrase every time before a connection can be created, and that is inconvenient.

To make working with passphrases a bit less complicated, the passphrase can be cached for a session. To do this, you need the **ssh-agent** and **ssh-add** commands. Assuming that the public/private key pair has already been created, this is an easy three-step procedure:

- Step 1.** Type **ssh-agent /bin/bash** to start the agent for the current (Bash) shell.
- Step 2.** Type **ssh-add** to add the passphrase for the current user's private key. The key is now cached.
- Step 3.** Connect to the remote server. Notice that there is no longer a need to enter the passphrase.

This procedure needs to be repeated for all new sessions that are created.

Summary

In this chapter, you learned how to configure the SSH server with advanced options. You also learned how to set security options for `sshd` and how to set specific client options that help in keeping connections alive for a longer period.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 20-3 lists a reference of these key topics and the page number on which each is found.

Table 20-3 Key Topics for Chapter 20

Key Topic Element	Description	Page
Table 20-2	Most useful sshd configuration options	451



Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

passphrase, connection

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. Which two commands do you need to cache the passphrase that is set on your private key?
2. You want to disallow root login and only allow user lisa to log in to your server. How would you do that?
3. How do you configure your SSH server to listen on two different ports?
4. What is the name of the main SSH configuration file?
5. When configuring a cache to store the passphrase for your key, where will this passphrase be stored?
6. What is the name of the file that contains SSH client settings for all users?
7. Which setting should you use to set the maximum number of concurrent SSH sessions to 10?
8. How do you configure SELinux to allow SSH to bind to port 2022?

9. How do you configure the firewall on the SSH server to allow incoming connections to port 2022?
10. Which setting could you use if you experience long timeouts while trying to establish an SSH connection?

End-of-Chapter Lab

In this end-of-chapter lab, you configure SSH for enhanced security and optimized connection settings. Use server1 to set up the SSH server, and use server2 as the SSH client.

Lab 20.1

1. Configure your SSH server in such a way that inactive sessions will be kept open for at least one hour.
2. Secure your SSH server so that it listens on port 2022 only and that only user lisa is allowed to log in.
3. Test the settings from server2. Make sure that the firewall as well as SELinux are configured to support your settings.

This page intentionally left blank



The following topics are covered in this chapter:

- Configuring a Basic Apache Server
- Understanding Apache Configuration Files
- Creating Apache Virtual Hosts

The following RHCSA exam objectives are covered in this chapter:

- No RHCSA exam objectives relate directly to Apache, but minimal Apache knowledge is required to master the SELinux-related objectives.

This is the only chapter in this book that discusses a subject that is not even listed in the RHCSA objectives. However, for a Red Hat server administrator, it is important to know how to deal with the Apache web service. In following chapters, you learn how to configure SELinux and installation servers. These are topics that are difficult to understand without knowing how to deal with the Apache web service. Also, in Chapter 11, “Working with Systemd,” you learned how to work with services in an RHEL 8 environment. Knowing how to configure a common service like the Apache web service will surely help doing so. That is why this chapter explains Apache web server basics.

Managing Apache HTTP Services

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 21-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 21-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Configuring a Basic Apache Server	1–4
Understanding Apache Configuration Files	5–7
Creating Apache Virtual Hosts	8–10

1. Which command installs the software packages that are needed to configure an Apache web server?
 - a. `yum install httpd`
 - b. `yum install web-server`
 - c. `yum install apache`
 - d. `yum install apache2`
2. What is the name of the main Apache configuration file?
 - a. `/etc/httpd/conf/httpd.conf`
 - b. `/etc/httpd/httpd.conf`
 - c. `/etc/apache2/apache.conf`
 - d. `/etc/httpd/default-server.conf`

3. Which parameter in the Apache configuration file is used to specify where Apache will serve its documents from?
 - a. **ServerRoot**
 - b. **ServerDocuments**
 - c. **DocumentRoot**
 - d. **DocumentIndex**
4. Which parameter in the main Apache configuration file defines the location where the Apache process looks for its configuration files?
 - a. **ServerRoot**
 - b. **ServerDocuments**
 - c. **DocumentRoot**
 - d. **DocumentIndex**
5. Which directory contains the main Apache configuration file?
 - a. /etc/httpd
 - b. /etc/httpd/conf
 - c. /etc/httpd/conf.d
 - d. /etc/httpd/conf.modules.d
6. Which directory contains the configuration files for the different Apache modules?
 - a. /etc/httpd
 - b. /etc/httpd/conf
 - c. /etc/httpd/conf.d
 - d. /etc/httpd/conf.modules.d
7. Which directory is used to drop configuration files that are installed from RPMs?
 - a. /etc/httpd
 - b. /etc/httpd/conf
 - c. /etc/httpd/conf.d
 - d. /etc/httpd/conf.modules.d

8. Which virtual host type allows you to run multiple virtual hosts on the same IP address?
 - a. Name-based
 - b. IP-based
 - c. Configuration-based
 - d. Default
9. Which line is used to start the definition of a virtual host that listens on port 80 of all IP addresses on the current server?
 - a. `<VirtualHost *:80>`
 - b. `<VirtualHost *>`
 - c. `<NameHost *:80`
 - d. `<NameHost *>`
10. Which of the following statements about virtual hosts is *not* true?
 - a. When virtual hosts are offered through an httpd process, the default configuration no longer works.
 - b. The names of virtual hosts must be resolvable through /etc/hosts or DNS.
 - c. To use virtual hosts, the mod_virt package must be installed.
 - d. Virtual host configurations can be specified in httpd.conf.

Foundation Topics

Configuring a Basic Apache Server

Configuring a basic Apache server is not hard to do. It consists of a few easy steps:

- Step 1.** Install the required software.
- Step 2.** Identify the main configuration file.
- Step 3.** Create some web server content.

Installing the Required Software

The Apache server is provided through some different software packages. The basic package is `httpd`; this package contains everything that is needed for an operational but basic web server. There are some additional packages, as well. For a complete overview, you can use the `yum search http` command, and you can use `yum module install httpd` to install the base package and common additional modules.

Identifying the Main Configuration File

The configuration of the Apache web server goes through different configuration files. The section “Understanding Apache Configuration Files” later in this chapter provides an overview of the way these files are organized. The main Apache configuration file is `/etc/httpd/conf/httpd.conf`. In this section, many parameters are specified. The most important parameter to understand for setting up a basic web server is the `DocumentRoot` parameter. This parameter specifies the default location where the Apache web server looks for its contents.

Another important configuration parameter is the `ServerRoot`. This defines the default directory where Apache will look for its configuration files. By default, the `/etc/httpd` directory is used for this purpose, but alternative directories can be used as well. Many other configuration files are referenced in the `httpd.conf` file, a portion of which is shown in Example 21-1. The use of additional configuration files makes it easy for applications to install snap-in files that will be included by the Apache server from RPM packages. The names of all these configuration files are relative to the `ServerRoot /etc/httpd`.

Example 21-1 Partial Contents of the /etc/httpd/conf/httpd.conf Configuration File

```
[root@localhost ~]# grep -v '#' /etc/httpd/conf/httpd.conf

ServerRoot "/etc/httpd"

Listen 80

Include conf.modules.d/*.conf

User apache
Group apache

ServerAdmin root@localhost

<Directory />
    AllowOverride none
    Require all denied
</Directory>

DocumentRoot "/var/www/html"

<Directory "/var/www">
    AllowOverride None
    Require all granted
</Directory>

<Directory "/var/www/html">
    Options Indexes FollowSymLinks

    AllowOverride None

    Require all granted
</Directory>

<IfModule dir_module>
    DirectoryIndex index.html
</IfModule>

<Files ".ht*">
    Require all denied
</Files>
```



```

ErrorLog "logs/error_log"

LogLevel warn

<IfModule log_config_module>
    LogFormat "%h %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%
        {User-Agent}i\"" combined
    LogFormat "%h %l %u %t \"%r\" %>s %b" common

    <IfModule logio_module>
        LogFormat "%h %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%
            {User-Agent}i\" %I %O" combinedio
    </IfModule>

    CustomLog "logs/access_log" combined
</IfModule>

<IfModule alias_module>

    ScriptAlias /cgi-bin/ "/var/www/cgi-bin/"

</IfModule>

<Directory "/var/www/cgi-bin">
    AllowOverride None
    Options None
    Require all granted
</Directory>

<IfModule mime_module>
    TypesConfig /etc/mime.types

    AddType application/x-compress .Z
    AddType application/x-gzip .gz .tgz

    AddType text/html .shtml
    AddOutputFilter INCLUDES .shtml
</IfModule>

```

```
AddDefaultCharset UTF-8

<IfModule mime_magic_module>
    MIMEMagicFile conf/magic
</IfModule>

EnableSendfile on
IncludeOptional conf.d/*.conf
```

Creating Web Server Content

After identifying the web server `DocumentRoot`, you know all you need to know to configure a basic web server. The Apache web server by default looks for a file with the name `index.html` and will present the contents of that document to clients using a browser to access the web server. It suffices to configure this file with very basic contents; just a line like “Welcome to my web server” will do.

To test the web server, you can launch a browser. The Firefox browser is installed by default on all graphical installations of RHEL 8. If your server does not run a graphical interface, use **yum install curl** to work with Apache from the command line.

In Exercise 21-1, you learn how to set up a basic Apache web server—nothing fancy, just enough to get you going and test web server functionality.

Exercise 21-1 Setting Up a Basic Web Server

1. Type **yum module install httpd**. This installs the `httpd` package and some of the most commonly used additional packages as well.
2. Open the main Apache configuration file with an editor, and look up the line that starts with `DocumentRoot`. This identifies the location where the Apache server will look for the contents it will service. Confirm that it is set to `/var/www/html`.
3. In the directory `/var/www/html`, create a file with the name **index.html**. In this file, type “Welcome to my web server”.
4. To start and enable the web server, type **systemctl enable --now httpd**. This starts the web server and makes sure that it starts automatically after restarting the server. Use **systemctl status httpd** to check that the web server is up and running. In Example 21-2 you can see what the result of this command should look like.
5. Type **yum install curl** to install the `elinks` text-based browser. Type **curl http://localhost** to connect to the web server and verify it is working.

Example 21-2 Verifying the Availability of the Apache Web Server with **systemctl status**

```
[root@localhost ~]# systemctl status httpd
httpd.service - The Apache HTTP Server
Loaded: loaded (/usr/lib/systemd/system/httpd.service; enabled;
       vendor preset: disabled)
Active: active (running) since Fri 2019-07-05 03:06:02 EDT; 2s ago
       Docs: man:httpd.service(8)
Main PID: 4540 (httpd)
Status: "Started, listening on: port 443, port 80"
Tasks: 213 (limit: 11222)
Memory: 24.2M
CGroup: /system.slice/httpd.service
        └─4540 /usr/sbin/httpd -DFOREGROUND
        └─4542 /usr/sbin/httpd -DFOREGROUND
        └─4543 /usr/sbin/httpd -DFOREGROUND
        └─4544 /usr/sbin/httpd -DFOREGROUND
        └─4545 /usr/sbin/httpd -DFOREGROUND

Jul 05 03:06:02 localhost.localdomain systemd[1]: Starting The Apache
HTTP Server...
Jul 05 03:06:02 localhost.localdomain httpd[4540]: AH00558: httpd:
Could not reliably determine the server's
Jul 05 03:06:02 localhost.localdomain httpd[4540]: Server configured,
listening on: port 443, port 80
Jul 05 03:06:02 localhost.localdomain systemd[1]: Started The Apache
HTTP Server.
```

Understanding Apache Configuration Files

A default installation of the Apache web server creates a relatively complex configuration tree in the `/etc/httpd` directory. Example 21-3 shows the default contents of this directory. The contents of this directory may differ on your server if additional software has been installed. Apache is modular, and upon installation of additional Apache modules, different configuration files might be installed here.

Example 21-3 Default Contents of the `/etc/httpd` Directory

```
[root@server1 httpd]# \ls -l
total 8
drwxr-xr-x. 2 root root 35 Feb 23 03:12 conf
drwxr-xr-x. 2 root root 4096 Feb 25 12:41 conf.d
drwxr-xr-x. 2 root root 4096 Feb 25 12:41 conf.modules.d
lrwxrwxrwx. 1 root root 19 Feb 17 13:26 logs -> ../../var/log/httpd
lrwxrwxrwx. 1 root root 29 Feb 17 13:26 modules -> ../../usr/lib64/
httpd/modules
lrwxrwxrwx. 1 root root 10 Feb 17 13:26 run -> /run/httpd
```

The first thing you notice is the presence of three symbolic links to logs, modules, and a run directory. These are created to allow Apache to be started in a chroot environment.

A chroot environment provides a fake root directory. This is a directory in the file system that is presented as the root directory for the process that is running in the chroot environment. This is done for security reasons: Processes that are running in a chroot environment can access files in that chroot environment only, which decreases the risk of security incidents occurring when intruders manage to get a login shell using the web server identity and try walking through the file system to do unauthorized things.

**Key
Topic**

The main configuration files for the Apache web server are in the `/etc/httpd/conf` directory. To start, there is the `httpd.conf` file, which contains the most important configuration parameters. Apart from that, there is a file with the name `magic`. This file is used by the browser to interpret how the contents of the web server should be displayed. It makes sure that the web server content is shown correctly in different browsers.

The `/etc/httpd/conf.d` directory contains files that are included in the Apache configuration. Adding files is done by the line `Include conf.d/*.conf` in the `httpd.conf` file. This directory can be used by RPMs that include Apache snap-in files. As is the case for the `ServerRoot`, this approach makes it possible to add configuration files that define the different web pages without changing the contents of the `/etc/httpd/conf/httpd.conf` file.

The last configuration directory is `/etc/httpd/conf.modules.d`. Apache is a modular web server. Therefore, the functionality of the Apache web server can easily be extended by adding additional modules that enable many different features. If modules are used, they can use their own module-specific configuration files, which will be dropped in the `/etc/httpd/conf.modules.d` directory. Again, the purpose of this approach is to keep the configuration in `/etc/httpd/conf.d/httpd.conf` as clean as possible and to make sure that module-specific configuration is not overwritten if the Apache generic configuration is updated.

Creating Apache Virtual Hosts

Many companies host more than one website. Fortunately, it is not necessary to install a new Apache server for every website that you want to run. Apache can be configured to work with virtual hosts. A virtual host is a distinct Apache configuration file or section that is created for a unique hostname. When working with virtual hosts, the procedure to access the host is roughly like the following:

1. The client starts a session to a specific virtual host, normally by starting a browser and entering the URL to the website the client wants to use.

2. DNS helps resolve the IP address of the virtual host, which is the IP address of the Apache server that can host different virtual hosts.
3. The Apache process receives requests for all the virtual hosts it is hosting.
4. The Apache process reads the HTTP header of each request to analyze which virtual host this request needs to be forwarded to.
5. Apache reads the specific virtual host configuration file to find which document root is used by this specific virtual host.
6. The request is forwarded to the appropriate contents file in that specific document root.

When working with virtual hosts, there are a few things to be aware of:

**Key
Topic**

- If your Apache server is configured for virtual hosts, all servers it is hosting should be handled by virtual hosts. To create a catch-all entry for all HTTP requests that are directed to this host but that do not have a specific virtual host file, you can create a virtual host for `_default_:80`. If you don't do that, packages that successfully arrive on your server via DNS name resolution but don't find a matching virtual host are sent to the virtual host whose configuration the Apache process finds first. That leads to unpredicted results.
- Name-based virtual hosting is the most common solution. In this solution, virtual hosts use different names but the same IP address.
- IP-based virtual hosts are less common but are required if the name of a web server must be resolved to a unique IP address. IP-based virtual hosts do require several IP addresses on the same machine and are common in configurations where the Apache server uses TLS to secure connections.

Configuring virtual hosts is not an RHCSA objective, but it is useful to know how to configure them as a Linux administrator. Therefore, Exercise 21-2 walks you through the procedure.

Exercise 21-2 Configuring Apache Virtual Hosts

In this exercise, you create two virtual hosts. To set up virtual hosts, you first set up name resolution, after which you create the virtual hosts' configuration. Because SELinux has not been discussed yet, you temporarily switch off SELinux.

NOTE I later tell you that you should never switch off SELinux. For once, I make an exception to this important security rule. To focus on what needs to be done on the Apache web server, it is easier to focus just on Apache and not to configure SELinux as well.

1. On both server1 and server2, open the file /etc/hosts with an editor and add two lines that make it possible to resolve the names of the virtual host you are going to create to the IP address of the virtual machine:

```
192.168.4.210 server1.example.com server1
192.168.4.220 server2.example.com server2
192.168.4.210 account.example.com account
192.168.4.210 sales.example.com sales
```

2. On server1, open a root shell and add the following to the /etc/httpd/conf/httpd.conf file. (You can leave all other settings as they are.)

```
<Directory/www/docs>
    Require all granted
    AllowOverride None
</Directory>
```

3. On server1, open a root shell and create a configuration file with the name **account.example.com.conf** in the directory /etc/httpd/conf.d. Give this file the following content:

```
<VirtualHost *:80>
    ServerAdmin webmaster@account.example.com
    DocumentRoot /www/docs/account.example.com
    ServerName account.example.com
    ErrorLog logs/account.example.com-error_log
    CustomLog logs/account.example.com-access_log common
</VirtualHost>
```

4. Close the configuration file, and from the root shell type **mkdir -p /www/docs/account.example.com**.
5. Create a file with the name **index.html** in the account document root, and make sure its contents read “Welcome to account.”
6. Temporarily switch off SELinux using **setenforce 0**.
7. Use **systemctl restart httpd** to restart the Apache web server.
8. Use **curl http://account.example.com**. You should now see the account welcome page. (You may have to install elinks, using **yum install -y curl**.)
9. Back on the root shell, copy the /etc/httpd/conf.d/account.example.com.conf file to a file with the name /etc/httpd/conf.d/sales.example.com.conf.
10. Open the sales.example.com.conf file in **vi**, and use the **vi** command **:%s/account/sales/g**. This should replace all instances of “account” with the “sales.”
11. Create the **/www/docs/sales.example.com** document root, and create a file **index.html** in it, containing the text “Welcome to the sales server.”
12. Restart httpd and verify that both the account and the sales servers are accessible.

Summary

In this chapter, you learned about Apache basics. The information in this chapter helps you configure a basic Apache web server, which helps testing advanced topics like firewall configuration or SELinux configuration that are covered in subsequent chapters.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 21-2 lists a reference of these key topics and the page number on which each is found.



Table 21-2 Key Topics for Chapter 21

Key Topic Element	Description	Page
Paragraph	chroot explanation	465
List	Virtual host explanation	466

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

DocumentRoot, virtual host, chroot

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. Which **yum** group contains many useful Apache packages?
2. How do you enable the `httpd` service to be started automatically when booting?
3. What is the default location where RPMs can drop plug-in configuration files that should be considered by the Apache server?
4. Which command enables you to test a web server from a server that does not offer a graphical interface?
5. What is the name of the default Apache configuration file?
6. Which directory is used as the default Apache document root?
7. Which file does the Apache process look for by default in the document root?
8. Which command enables you to see whether the Apache web server is currently running?
9. Which location is preferable for storing virtual host configuration files?
10. Names of configuration files and directories in the main Apache configuration file are relative to the `ServerRoot`. To which directory is the `ServerRoot` set by default?

End-of-Chapter Lab

In this end-of-chapter lab, you install and configure a basic Apache web server.

Lab 21.1

1. Install the required packages that allow you to run a basic web server. Make sure that the web server process is started automatically when your server reboots. Do *not* use a virtual server.
2. Use **curl** to make sure the web server presents a default page showing “Welcome to my web server.”
3. Type **yum install httpd-manual** to install the Apache documentation.
4. Use a browser to test access to the `/manual` web page on your server.



The following topics are covered in this chapter:

- Understanding SELinux Working Modes
- Understanding Context Settings and the Policy
- Restoring Default File Contexts
- Using Boolean Settings to Modify SELinux Settings
- Diagnosing and Addressing SELinux Policy Violations

The following RHCSA exam objectives are covered in this chapter:

- Set enforcing and permissive modes for SELinux
- List and identify SELinux file and process context
- Restore default file contexts
- Use boolean settings to modify system SELinux settings
- Diagnose and address routine SELinux policy violations

Managing SELinux

Since the earliest days of Linux, file permissions have been the standard method of securing Linux systems. In some cases, file permissions are just not enough to secure a server fully. Let's take a look at an example:

One morning I found out that my server was hacked. An intruder had broken through a bad script on my web server and had obtained shell access as the `httpd` user—this was possible due to a bug in the shell code that I was using. Using this file access, he managed to create thousands of little PHP scripts that were involved in a massive DDoS attack.

From a security perspective, it is interesting that nothing really was wrong with the security settings on this server. All permissions were set in a decent way, and the `httpd` user, like any other user on a Linux server, does have permissions to create files in `/var/tmp`, as in `/tmp`. So, what would have been a good solution to prevent this kind of problem?

You could, of course, argue that the administrator of the web server should have been doing a better job and should have been watching what the scripts on the server were doing. But that is not how Linux servers are normally used. The Linux server administrator does not necessarily have in-depth knowledge of the internals of all the applications running on the Linux server, and the application administrator does not understand enough about Linux to ensure that something like this can never happen.

Another solution is to apply further security measures. For instance, this specific situation would have been prevented if the permission to run program files from the `/tmp` and `/var/tmp` directory had been taken away by using the **`noexec mount`** option. But even if that would have worked for this specific situation, it is not a good overall security solution that prevents applications from doing things they are not supposed to be doing. Basically, Linux just needs a default security solution that covers all settings.

That is why SELinux was invented. SELinux provides mandatory access control to a Linux server, where every system call is denied unless it has been specifically allowed. This chapter explains how to use SELinux to make sure that serious security incidents will never happen on your server.

TIP By any means, make sure that at the end of the exam SELinux is working on your server. If it is not, it will cost you many points!

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 22-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 22-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Understanding SELinux Working Modes	1–2
Understanding Context Settings and Policy	3–6
Restoring Default File Contexts	7
Using Boolean Settings to Modify SELinux Settings	8
Diagnosing and Addressing SELinux Policy Violations	9–10

- Which of the following is not a valid SELinux mode?
 - Enforcing
 - Permissive
 - Disabled
 - Enabled
- Which of the following commands enable you to see the current SELinux mode? (Choose two.)
 - sestatus
 - lsmode
 - semode
 - getenforce

3. Which of the following items in the context label is the most significant for SELinux system administration tasks?
 - a. Type
 - b. User
 - c. Role
 - d. Mission
4. To which of the following can SELinux security *not* be applied?
 - a. Users.
 - b. Files.
 - c. Ports.
 - d. It can be applied to all of the above.
5. Which command-line switch is used with many commands to display SELinux-related information?
 - a. **-S**
 - b. **-X**
 - c. **-Z**
 - d. **-D**
6. Which of the following commands should be used to set the context type of the directory /web to httpd_sys_content_t?
 - a. **chcon -t httpd_sys_content_t /web**
 - b. **semanage -t httpd_sys_content_t "/web(/.*)?"**
 - c. **semanage fcontext -t httpd_sys_content_t "/web(/.*)?"**
 - d. **semanage fcontext -a -t httpd_sys_content_t "/web(/.*)?"**
7. Which command must you run to ensure that a file has the appropriate SELinux context after moving the file to another location?
 - a. **reboot**
 - b. **restorecon /new/filename**
 - c. **chcon**
 - d. **restorecon -R /etc/selinux -v**

8. Which command enables you to change a Boolean in a way that it survives a reboot?
- a. **chcon boolean -P**
 - b. **setsebool -P**
 - c. **setsebool**
 - d. **semanage boolean**
9. Which file contains all information you need to troubleshoot SELinux messages?
- a. **/var/log/audit/audit.log**
 - b. **/var/log/selinux/selinux.log**
 - c. **/var/log/messages**
 - d. **/var/log/selinux.log**
10. You want to grep the log file for SELinux log messages. Which of the following strings should you use **grep** on?
- a. **selinux**
 - b. **deny**
 - c. **violation**
 - d. **avc**

Foundation Topics

Understanding SELinux Working and Modes

If SELinux is enabled and nothing else has been configured, all system calls are denied. To specify what exactly is allowed, a policy is used. In this policy, rules define which source domain is allowed to access which target domain. The source domain is the object that is trying to access something. Typically, this is a process or a user. The target domain is the object that is accessed. Typically, that is a file, a directory, or a network port. To define exactly what is allowed, context labels are used. Using these labels is the essence of SELinux because these labels are used to define access rules. Table 22-2 summarizes the most important SELinux building blocks.

Table 22-2 SELinux Core Elements

Element	Use
Policy	A collection of rules that define which source has access to which target.
Source domain	The object that is trying to access a target. Typically a user or a process.
Target domain	The thing that a source domain is trying to access. Typically a file or a port.
Context	A security label that is used to categorize objects in SELinux.
Rule	A specific part of the policy that determines which source domain has which access permissions to which target domain.
Labels	Same as a context label, defined to determine which source domain has access to which target domain.

Key
Topic

On a Linux system, you can choose to enable or disable SELinux. When SELinux is enabled, kernel support for SELinux is loaded, and some applications that are SELinux aware change their behavior, because specific libraries are used on a system that has SELinux enabled. If SELinux is disabled, no SELinux activity happens at all. Changing between SELinux enabled mode and SELinux disabled mode requires a reboot of your system. This is because SELinux is a feature that is deeply interwoven with the Linux kernel.

If on a system SELinux is enabled, you can select to put SELinux in enforcing mode or in permissive mode. In enforcing mode, SELinux is fully operational and enforcing all SELinux rules in the policy. If SELinux is in permissive mode, all SELinux-related activity is logged, but no access is blocked. This makes SELinux

permissive mode an excellent mode to do troubleshooting. Permissive mode is also a great way to do something and see the result from an SELinux perspective by analyzing the messages that are written to `/var/log/audit/audit/log`. That can help in building new and more efficient policies.

To set the default SELinux mode while booting, use the file `/etc/sysconfig/selinux`. Example 22-1 shows the content of this file.

Example 22-1 Content of the `/etc/sysconfig/selinux` File

```
[root@server1 ~]# cat /etc/sysconfig/selinux

# This file controls the state of SELinux on the system.
# SELINUX= can take one of these three values:
#     enforcing - SELinux security policy is enforced.
#     permissive - SELinux prints warnings instead of enforcing.
# disabled - No SELinux policy is loaded.
SELINUX=enforcing
# SELINUXTYPE= can take one of these two values:
#     targeted - Targeted processes are protected,
#     minimum - Modification of targeted policy. Only selected
#               processes are protected.
#     mls - Multi Level Security protection.
SELINUXTYPE=targeted
```

As you can see, in this file, which is read while booting, you can choose to put SELinux in enforcing, permissive, or disabled mode.

On a server that currently has SELinux enabled, you can use the **getenforce** command to see whether SELinux currently is in enforcing mode or in permissive mode. To switch between permissive mode and enforcing mode, you can use **setenforce**. The command **setenforce 0** temporarily puts SELinux in permissive mode, and **setenforce 1** puts SELinux temporarily in enforcing mode. To change the default mode persistently, you need to write it to `/etc/sysconfig/selinux`.

Another useful command is **sestatus**. If used with the option **-v**, this command shows detailed information about the current status of SELinux on a server. Example 22-2 shows the output of the **sestatus -v** command. It not only shows you which parts of SELinux are enabled but also shows the current version of the policy that is loaded and the context labels for some critical parts of the system.

Example 22-2 Using **sestatus -v** to Get Detailed Information About the Current Protection Status

```
[root@server1 ~]# sestatus -v
SELinux status:                enabled
SELinuxfs mount:              /sys/fs/selinux
SELinux root directory:      /etc/selinux
Loaded policy name:           targeted
Current mode:                 enforcing
Mode from config file:       enforcing
Policy MLS status:           enabled
Policy deny_unknown status:   allowed
Memory protection checking:   actual (secure)
Max kernel policy version:    31

Process contexts:
Current context:              unconfined_u:unconfined_r:unconfined_
                             t:s0-s0:c0.c1023
Init context:                 system_u:system_r:init_t:s0
/usr/sbin/sshd                system_u:system_r:sshd_t:s0-s0:c0.c1023

File contexts:
Controlling terminal:         unconfined_u:object_r:user_devpts_t:s0
/etc/passwd                   system_u:object_r:passwd_file_t:s0
/etc/shadow                   system_u:object_r:shadow_t:s0
/bin/bash                     system_u:object_r:shell_exec_t:s0
/bin/login                    system_u:object_r:login_exec_t:s0
/bin/sh                       system_u:object_r:bin_t:s0 ->
                             system_u:object_r:shell_exec_t:s0
/sbin/agetty                  system_u:object_r:getty_exec_t:s0
/sbin/init                    system_u:object_r:bin_t:s0 ->
                             system_u:object_r:init_exec_t:s0
/usr/sbin/sshd                system_u:object_r:sshd_exec_t:s0
```

In Exercise 22-1, you practice working with these different modes.

Exercise 22-1 Manipulating SELinux Modes

1. Open a root console on your server and type **getenforce**. You'll normally see that SELinux is in enforcing mode. If the output of **getenforce** shows SELINUX=disabled, edit the /etc/sysconfig/selinux file and change SELINUX= to **enforcing** (that is, SELINUX=**enforcing**). Then save the file and reboot the system before you continue.
2. Type **setenforce 0** and type **getenforce** again. SELinux now switches to permissive mode.
3. Open the file /etc/sysconfig/selinux with an editor and change the line SELINUX= so that it reads SELINUX=**disabled**. Reboot your server.
4. After rebooting, log in to a root shell again and type **getenforce**. You'll see that SELinux is now in disabled mode.
5. Try using the command **setenforce 1**. You'll see the message "setenforce: SELinux is disabled." You cannot switch between disabled and enforcing mode without rebooting your server.
6. Open the file /etc/sysconfig/selinux again and change the line SELINUX=disabled back to SELINUX=**enforcing**. Reboot your system again.
7. After rebooting, type **sestatus -v** and read current status information about SELinux.

TIP Whatever you do, do not change the contents of the /etc/sysconfig/selinux file on the exam. Your exam system must be configured with SELinux in enforcing mode. To troubleshoot SELinux, you can put it temporarily in permissive mode by using **setenforce 0**. In this mode, you can troubleshoot any SELinux problem, but at least you'll be sure that after a reboot your server is started in a mode where SELinux is enabled, which is an essential requirement if you want to pass the exam.

Note that on real Red Hat servers, SELinux on occasion is set to be disabled. Putting SELinux in disabled mode certainly makes it easier for administrators to run their applications. However, it also makes the server much less secure. Often, ignorance of the system administrator is the only reason SELinux is put in disabled mode. If an application vendor tells you that the application is supported only if SELinux is disabled, that simply means the application vendor has no knowledge about SELinux.

A fully enforcing system is especially important if your server is accessed directly by users from the Internet. If your server cannot be reached directly from the Internet and is in a safe internal network, having SELinux enabled is not strictly necessary (but I recommend always keeping it in enforcing mode anyway). On the RHCSA

exam, however, you must make sure that SELinux is enabled and fully protecting your server.

NOTE SELinux is often disabled on servers either because of administrator laziness or because an application vendor who doesn't know how to deal with it has instructed the administrator to disable it. On many occasions, even applications that do not know how to work with SELinux can be fully functional on a server with SELinux. It just takes a bit more work to figure out the additional rules in the policy that need to be created to use the application on an SELinux-enabled system.

Understanding Context Settings and the Policy

Context settings are an important part of SELinux operations. The context is a label that can be applied to different objects:

- Files and directories
- Ports
- Processes
- Users

**Key
Topic**

Context labels define the nature of the object, and SELinux rules are created to match context labels of source objects to the context labels of target objects. So, setting correct context labels is a very important skill for system administrators. You learn how to do that later in this chapter.

NOTE Managing SELinux context labels is a key skill for securing systems with SELinux. It is not listed in the RHCSA exam objectives though. Nevertheless, I'll give you a decent explanation of how context labels work, because a mismatch of context labels can create lots of SELinux-related problems.

Monitoring Current Context Labels

To see current context settings on the objects in the previous bulleted list, many commands offer support for the **-Z** option. In Example 22-3, you see how **ls -Z** shows context settings for some directories in the **/** file system. Other commands also support the **-Z** option to show current context label settings. Some examples are **ps -Zaux**, which shows a list of all processes, including their context label, or **ss -Ztul**, which shows all network ports and the current context label associated with each port.

Example 22-3 Displaying Context Labels on Files with **ls -Z**

```
[root@server1 /]# ls -Z
system_u:object_r:bin_t:s0 bin      unconfined_u:object_r:
                                default_t:s0 repo
system_u:object_r:boot_t:s0 boot   system_u:object_r:admin_
                                home_t:s0 root
system_u:object_r:device_t:s0 dev   system_u:object_r:var_
                                run_t:s0 run
system_u:object_r:etc_t:s0 etc      system_u:object_r:bin_t:
                                s0 sbin
system_u:object_r:unlabeled_t:s0 files system_u:object_r:var_t:
                                s0 srv
system_u:object_r:home_root_t:s0 home system_u:object_r:root_t:
                                s0 stratis
system_u:object_r:lib_t:s0 lib      system_u:object_r:unlabeled_
                                t:s0 stratis1
system_u:object_r:lib_t:s0 lib64    system_u:object_r:sysfs_t:
                                s0 sys
system_u:object_r:mnt_t:s0 media    system_u:object_r:tmp_t:
                                s0 tmp
system_u:object_r:mnt_t:s0 mnt      system_u:object_r:usr_t:
                                s0 usr
system_u:object_r:usr_t:s0 opt      system_u:object_r:var_t:
                                s0 var
system_u:object_r:proc_t:s0 proc    system_u:object_r:
                                unlabeled_t:s0 vdo1
```

Every context label always consists of three different parts:

**Key
Topic**

- **User:** The user can be recognized by `_u` in the context label; it is set to `system_u` on most directories in Example 22-3. SELinux users are not the same as Linux users, and they are not important on the RHCSA exam.
- **Role:** The role can be recognized by `_r` in the context label. In Example 22-3, most objects are labeled with the `object_r` role. In advanced SELinux management, specific SELinux users can be assigned permissions to specific SELinux roles. For the RHCSA exam, you do not have to know how to configure roles.
- **Type:** The type context can be recognized by `_t` in the context label. In Example 22-3, you can see that a wide variety of context types are applied to the directories in the `/` file system. Make sure that you know how to work with context types, because they are what the RHCSA exam expects you to know.

TIP Just to make sure that you are focusing on the parts that really matter on the RHCSA exam, you need to work with context types only. You can safely ignore the user and role parts of the context label.

Setting Context Types

As an administrator, it is important that you know how to set context types. You can set these context types on files and directories and other objects such as network ports. Let's focus on that task first.

You can use two commands to set context type:

- **semanage:** This is the command you want to use. The **semanage** command writes the new context to the SELinux policy, from which it is applied to the file system.
- **chcon:** This command is for use in specific cases only and normally should be avoided. The **chcon** command writes the new context to the file system and not to the policy. Everything that is applied with **chcon** is overwritten when the file system is relabeled, or the original context is restored from the policy to the file system. Do *not* use this command!

NOTE You might want to know why I bother mentioning **chcon** if you should not use it. Well, you'll often see the **chcon** command still being referred to in the documentation, which might give the impression that it is a useful command. It is not, because if your file system is relabeled, all changes applied with **chcon** are lost. File system relabeling actions can take you by surprise if you are new to SELinux, and you will fail your exam if by accident a file system relabeling happens on a file system where you have applied SELinux context with **chcon**. So, I repeat: Do *not* use it.

TIP The **semanage** command may not be installed by default. Fortunately, you can type **yum whatprovides */semanage** to find the RPM containing **semanage** and then install it. Do not learn the names of all relevant RPMs by heart; just remember **yum whatprovides**. It will find any RPM you need. See Chapter 9, "Managing Software," for more information about the use of the **yum** command and package management in general.

To set context using **semanage**, you first need to find the appropriate context (a topic covered in more depth in the next section, “Finding the Context Type You Need”). An easy way to find the appropriate context is by looking at the default context settings on already-existing items. If you want to change the context for a web server, for example, type **ls -Z /var/www** to see the context settings:

```
[root@server1 ~]# ls -Z /var/www
drwxr-xr-x. root root system_u:object_r:httpd_sys_script_exec_t:s0
cgi-bin
drwxr-xr-x. root root system_u:object_r:httpd_sys_content_t:s0 html
```

As you can see, the context settings on `/var/www/html` are set to `httpd_sys_content_t`. (As a reminder, we’re looking only at the context type because the user and role are for advanced use only.) To set this context type to any new directory that you want to be accessible by the Apache web server, use the following command:

```
semanage fcontext -a -t httpd_sys_content_t "/mydir(/.*)?"
```

In this command, the option **-a** is used to add a context type. This is what you need to do for all directories that you have created manually yourself. Then you use **-t** to change the context type (as opposed to user and role). The last part of the command is a regular expression, which is used to refer to the directory `/mydir` and anything that might exist below this directory.

Setting the context in this way is not enough, though, because you’ll write it only to the policy and not to the file system. To complete the command, you need to apply the policy setting to the file system, as follows:

```
restorecon -R -v /mydir
```

You’ll see that the new context is now applied, which allows the `httpd` process to access the directory.

TIP The **semanage** command is not the easiest command to remember. Fortunately, it has some excellent man pages. Type **man semanage** and use **G** to go all the way down to the bottom of the man page. You’ll now see the “See Also” section, which mentions **semanage-fcontext**, which is about managing file context with **semanage**. Open this man page using **man semanage-fcontext**, type **/examples**, and you’ll see some pretty examples that mention exactly what you need to know (see Example 22-4).

Example 22-4 semanage fcontext Usage Example from the man Page**EXAMPLE**

```
remember to run restorecon after you set the file context
Add file-context for everything under /web
# semanage fcontext -a -t httpd_sys_content_t "/web(/.*)?"
# restorecon -R -v /web
```

Substitute /home1 with /home when setting file context

```
# semanage fcontext -a -e /home /home1
# restorecon -R -v /home1
```

For home directories under top level directory, for example /disk6/home,

execute the following commands.

```
# semanage fcontext -a -t home_root_t "/disk6"
# semanage fcontext -a -e /home /disk6/home
# restorecon -R -v /disk6
```

SEE ALSO

selinux (8), semanage (8)

AUTHOR

This man page was written by Daniel Walsh <dwalsh@redhat.com>

20130617

semanage-fcontext(8)

Now it is time for an exercise. In Exercise 22-2, you learn how to change the document root for the Apache web server and label the new document root in the right way.

Exercise 22-2 Setting a Context Label on a Nondefault Apache Document Root

1. Open a root shell and type **yum install httpd curl -y**.
2. Still from the root shell, type **mkdir /web**.
3. Type **vim /web/index.html** and put the following contents in the file: **welcome to my web server**.
4. Type **vim /etc/httpd/conf/httpd.conf** to open the Apache configuration file and find the **DocumentRoot** parameter. Change it so that it reads **DocumentRoot "/web"**.

5. In the same `httpd.conf` configuration file, add the following section, as without this section it will be Apache and not SELinux blocking access to the new DocumentRoot:

```
<Directory "/web">
    AllowOverride None
    Require all granted
</Directory>
```
6. Type **`systemctl enable --now httpd`** to start and enable the httpd service. Note that if the httpd service was already running, you'll need to use **`systemctl restart httpd`** to restart it so that it can pick up the changes you've made to the `httpd.conf` configuration file.
7. Type **`curl http://localhost`**. You'll see the default Red Hat web page and not the contents of the `index.html` file you have just created.
8. Type **`setenforce 0`** to switch SELinux to permissive mode.
9. Repeat step 7. You'll now get access to your custom web page, which proves that SELinux was doing something to block access.
10. Type **`semanage fcontext -a -t httpd_sys_content_t "/web(/.*)"`** to apply the new context label to `/web`.
11. Type **`restorecon -R -v /web`**. The `-v` (verbose) option ensures that you see what is happening and that you will see the new context being applied to `/web`.
12. Set SELinux back in enforcing mode, using **`setenforce 1`**.
13. Type **`elinks http://localhost`**. You'll now get access to your custom web page because SELinux now allows access to it.

Finding the Context Type You Need

One of the challenging parts of setting SELinux contexts is finding the context you need. Roughly, there are three approaches:

- Look at the default environment.
- Read the configuration files.
- Use **`man -k _selinux`** to find SELinux-specific man pages for your service.

The most powerful way of getting the SELinux information you need is by using **`man -k _selinux`**, which searches the database of man pages for man pages that match `_selinux` in the name or description of the man page. On RHEL 8, however,

these man pages are not installed by default. To install them, you need to install the `policycoreutils-devel` package, after which you can use the command **sepolicy manpage -a -p /usr/share/man/man8** to install the SELinux man pages.

Exercise 22-3 guides you through the procedure you need to apply to install the application-specific SELinux man pages.

TIP Exercise 22-3 shows an essential skill. Make sure that you master this procedure before taking the exam.

Exercise 22-3 Installing SELinux-Specific Man Pages

1. Type **man -k _selinux**. You'll probably see just one or two man pages.
2. Type **yum provides */sepolicy**. This shows you the name of the RPM that contains the `sepolicy` binary, which is `policycoreutils-devel`.
3. Type **yum -y install policycoreutils-devel** to install this package.
4. Type **sepolicy manpage -a -p /usr/share/man/man8** to install the man pages.
5. Type **man -k _selinux**. You'll see no changes yet.
6. Type **mandb** to update the database that contains names and descriptions of all man pages that are installed.
7. Once the **mandb** command has finished (this can take a few minutes), type **man -k _selinux**. You'll now see a long list of man pages scrolling by.
8. Type **man -k _selinux | grep http** to find the man page that documents SELinux settings for the `httpd` service and scroll through it. Notice that it is a complete list of all that you can do with SELinux on the `httpd` service.

Restoring Default File Contexts

In the previous section, you learned how to apply context types using **semanage**. You also applied the context settings from the policy to the file system using **restorecon**. The **restorecon** command is a useful command because in the policy the default settings are defined for most files and directories on your computer. If the wrong context setting is ever applied, you just have to type **restorecon** to reapply it from the policy to the file system.

Using **restorecon** this way can be useful to fix problems on new files. Before explaining how to do it, let's take a look at how new context settings are applied:

Key Topic

- If a new file is created, it inherits the context settings from the parent directory.
- If a file is copied to a directory, this is considered a new file, so it inherits the context settings from the parent directory.
- If a file is moved, or copied while keeping its properties (by using **cp -a**), the original context settings of the file are applied.

Especially the latter of these three situations is easily fixed by using **restorecon**. Exercise 22-4 simulates this problem, and you fix it using **restorecon**.

It is also possible to relabel the entire file system. Doing so applies all context settings as defined in the policy to the file system. Because the policy should always be leading and contain correct context settings, relabeling a file system may be a good idea. To relabel the file system, you can either use the command **restorecon -Rv /** or create a file with the name **/.autorelabel**. The next time your server is restarted, the file system will automatically be relabeled.

A relabeling action sometimes occurs spontaneously. If while troubleshooting a server you have started the server in a mode where SELinux is disabled, and you have applied modifications to the file system, SELinux will detect that the file system has changed without SELinux monitoring it. This will result in an automatic relabeling of the entire file system. Note that on a large file system, relabeling the file system can take a significant amount of time.

Exercise 22-4 Using restorecon to Relabel Files

1. From a root shell, type **ls -Z /etc/hosts**. You'll see the file has the `net_conf_t` context label.
2. Type **cp /etc/hosts ~** to copy the file to the root home directory. Because copying is considered the creation of a new file, the context setting on the `~/hosts` file is set as `admin_home_t`. Use **ls -Z ~/hosts** to verify this.
3. Type **mv ~/hosts /etc** and confirm that you want to overwrite the existing file.
4. Type **ls -Z /etc/hosts** to confirm that the context type is still set to `admin_home_t`.
5. Type **restorecon -v /etc/hosts** to reapply the correct context type. The **-v** option shows you what is happening.
6. Type **touch /.autorelabel** and restart your server. While restarting, make sure to press the **Escape** key on your keyboard so that you'll see boot messages. You'll see that the file system is automatically relabeled.

Using Boolean Settings to Modify SELinux Settings

In the SELinux policy, there are many rules. Some of these rules allow specific activity, whereas other rules deny that activity. Changing rules is not easy, and that is why SELinux Booleans are provided to easily change the behavior of a rule.

An example of a Boolean is `ftpd_anon_write`, which by default is set to off. That means that even if you have configured your FTP server to allow anonymous writes, the Boolean will still deny it, and the anonymous user cannot upload any files. If a conflict exists between the setting of a parameter in a service configuration file and in a Boolean, the Boolean always takes precedence. But Booleans are easy to change.

To get a list of Booleans on your system, type **getsebool -a**. If you are looking for Booleans that are set for a specific service, use **grep** to filter down the results. In Example 22-5, you can see how this command is used to show current Booleans that match FTP.

An alternative way to show current Boolean settings is by using the **semanage boolean -l** command. This command provides some more details, because it shows the current Boolean setting and the default Boolean setting.

Example 22-5 Displaying Boolean Settings

```
root@server1 ~]# getsebool -a | grep ftp
ftp_home_dir --> off
ftpd_anon_write --> off
ftpd_connect_all_unreserved --> off
ftpd_connect_db --> off
ftpd_full_access --> off
ftpd_use_cifs --> off
ftpd_use_fusefs --> off
ftpd_use_nfs --> off
ftpd_use_passive_mode --> off
httpd_can_connect_ftp --> off
httpd_enable_ftp_server --> off
sftpd_anon_write --> off
sftpd_enable_homedirs --> off
sftpd_full_access --> off
sftpd_write_ssh_home --> off
tftp_anon_write --> off
tftp_home_dir --> off
```

To change a Boolean, you can use **setsebool**. If you want to switch the `ftpd_anon_write` Boolean to allow anonymous writes, for example, use **setsebool ftpd_anon_write on**. This changes the runtime value of the Boolean but does not change it permanently. To apply permanent changes to a Boolean, use **setsebool -P**. Notice that this takes longer, because parts of the policy need to be recompiled to apply the modification. In Exercise 22-5, you apply these commands to see how Booleans are working.

Exercise 22-5 Working with SELinux Booleans

1. From a root shell, type **getsebool -a | grep ftp**. You'll see the `ftpd_anon_write` Boolean, with its current value set to off.
2. Type **setsebool ftpd_anon_write on**. This changes the value in the runtime.
3. Type **getsebool ftpd_anon_write**. It shows the value of the Boolean as on.
4. Type **semanage boolean -l | grep ftpd_anon**. Notice that this command shows the runtime configuration set to on, but the permanent setting is still set to off.
5. Use **setsebool -P ftpd_anon_write on** to switch the runtime and the default setting for the Boolean to on.
6. Repeat **semanage boolean -l | grep ftpd_anon**. Notice that it is now set to on, on.

Diagnosing and Addressing SELinux Policy Violations

Configuring a system with SELinux can be a challenging task. To make it easier to understand what is happening, SELinux logs everything it is doing. The primary source to get logging information is the audit log, which is in `/var/log/audit/audit.log`. SELinux messages are logged with type=AVC in the audit log. So, to see what SELinux is doing, you can use the command **grep AVC /var/log/audit/audit.log**. If SELinux messages have been logged, this command shows a result as in Example 22-6.

Example 22-6 Getting SELinux Messages from audit.log

```
[root@server1 ~]# grep AVC /var/log/audit/audit.log | grep http
type=AVC msg=audit(1559986797.093:185): avc: denied { getattr }
  for pid=32939 comm="httpd" path="/web/index.html" dev="dm-0"
  ino=35321780 scontext=system_u:system_r:httpd_t:s0 tcontext=
  unconfined_u:object_r:default_t:s0 tclass=file permissive=0
type=AVC msg=audit(1559986797.093:186): avc: denied { getattr }
  for pid=32939 comm="httpd" path="/web/index.html" dev="dm-0"
  ino=35321780 scontext=system_u:system_r:httpd_t:s0 tcontext=
  unconfined_u:object_r:default_t:s0 tclass=file permissive=0
type=AVC msg=audit(1559986815.360:188): avc: denied { getattr }
  for pid=32939 comm="httpd" path="/web/index.html" dev="dm-0"
  ino=35321780 scontext=system_u:system_r:httpd_t:s0 tcontext=
  unconfined_u:object_r:default_t:s0 tclass=file permissive=0
type=AVC msg=audit(1559986815.360:189): avc: denied { getattr }
  for pid=32939 comm="httpd" path="/web/index.html" dev="dm-0"
  ino=35321780 scontext=system_u:system_r:httpd_t:s0 tcontext=
  unconfined_u:object_r:default_t:s0 tclass=file permissive=0
type=AVC msg=audit(1559986883.549:192): avc: denied { getattr }
  for pid=33214 comm="httpd" path="/web/index.html" dev="dm-0"
  ino=35321780 scontext=system_u:system_r:httpd_t:s0 tcontext=
  unconfined_u:object_r:default_t:s0 tclass=file permissive=0
type=AVC msg=audit(1559986883.550:193): avc: denied { getattr }
  for pid=33214 comm="httpd" path="/web/index.html" dev="dm-0"
  ino=35321780 scontext=system_u:system_r:httpd_t:s0 tcontext=
  unconfined_u:object_r:default_t:s0 tclass=file permissive=0
type=AVC msg=audit(1559986927.406:197): avc: denied { getattr }
  for pid=33214 comm="httpd" path="/web/index.html" dev="dm-0"
  ino=35321780 scontext=system_u:system_r:httpd_t:s0 tcontext=
  unconfined_u:object_r:default_t:s0 tclass=file permissive=1
type=AVC msg=audit(1559986927.406:198): avc: denied { read } for
  pid=33214 comm="httpd" name="index.html" dev="dm-0" ino=35321780 sco
  ntext=system_u:system_r:httpd_t:s0 tcontext=unconfined_u:object_r:
  default_t:s0 tclass=file permissive=1
type=AVC msg=audit(1559986927.406:198): avc: denied { open }
  for pid=33214 comm="httpd" path="/web/index.html" dev="dm-0"
  ino=35321780 scontext=system_u:system_r:httpd_t:s0 tcontext=
  unconfined_u:object_r:default_t:s0 tclass=file permissive=1
type=AVC msg=audit(1559986927.406:199): avc: denied { map }
  for pid=33214 comm="httpd" path="/web/index.html" dev="dm-0"
  ino=35321780 scontext=system_u:system_r:httpd_t:s0 tcontext=
  unconfined_u:object_r:default_t:s0 tclass=file permissive=1
```

At first sight, the SELinux log messages look complicated. If you look a bit closer, though, they are not that hard to understand. Let's take a closer look at the last line in the log file:

```
type=AVC msg=audit(1559986927.406:199): avc: denied { map } for
pid=33214 comm="httpd" path="/web/index.html" dev="dm-0"
ino=35321780 scontext=system_u:system_r:httpd_t:s0 tcontext=
unconfined_u:object_r:default_t:s0 tclass=file permissive=1
```

The first relevant part in this line is the text `avc: denied { map }`. That means that a map request was denied, so some process has tried to read attributes of a file and that was denied, because it is a policy violation. Following that message, we can see `comm=httpd`, which means that the command trying to issue the `getattr` request was `httpd`, and we can see `path="/web/index.html"`, which is the file that this process has tried to access.

In the last part of the log line, we can get information about the source context and the target context. The source context (which is the context setting of the **httpd** command) is set to `http_t`, and the target context (which is the context setting of the `/web/index.html` file) is set to `default_t`. And apparently, SELinux did not like that too much. So, to fix this, you would have to relabel the file, as discussed earlier in the chapter.

Making SELinux Analyzing Easier

Based on the information you find in the `audit.log`, you may be able to decide what you need to do to fix the problem. Because the information in the `audit.log` is not easy to understand, there is **sealert**. First, you may need to install **sealert** by using **yum -y install setroubleshoot-server**. Then, it is a good idea to restart your server to make sure that all processes that are involved are restarted correctly. The next time an SELinux message is written to the audit log, an easier-to-understand message is written to `syslog` and, by default, can be read in `/var/log/messages`. Example 22-7 shows an output example.

Example 22-7 sealert Makes Analyzing SELinux Logs Easier

```
Nov 2 10:01:40 server1 setroubleshoot: Plugin Exception restorecon
Nov 2 10:01:40 server1 setroubleshoot: SELinux is preventing /usr/
sbin/httpd from getattr access on the file. For complete SELinux
messages. run sealert -l 0ed02423-1149-4561-b6a0-8ea2957329ea
Nov 2 10:01:40 server1 python: SELinux is preventing /usr/sbin/httpd
from getattr access on the file.

***** Plugin catchall_labels (83.8 confidence) suggests
*****
```

```

If you want to allow httpd to have getattr access on the file
Then you need to change the label on $FIX_TARGET_PATH
Do
# semanage fcontext -a -t FILE_TYPE '$FIX_TARGET_PATH'
where FILE_TYPE is one of the following: NetworkManager_exec_t,
  NetworkManager_log_t,
--removed 5 pages of the "one of the following" output --
Then execute:
restorecon -v '$FIX_TARGET_PATH'

**** Plugin catchall (17.1 confidence) suggests
*****

If you believe that httpd should be allowed getattr access on the
file by default.
Then you should report this as a bug.
You can generate a local policy module to allow this access.
Do
allow this access for now by executing:
# grep httpd /var/log/audit/audit.log | audit2allow -M mypol
# semodule -i mypol.pp

```

The useful thing about **sealert** is that it tries to analyze what has happened and, based on the analysis, suggests what you need to do to fix the problem. The not-so-useful part is that in some cases (as was the case in this example), hundreds of possible context types are shown, and the administrator has to choose the right one. (I removed five pages of output in Example 22-6 to keep it readable.) So, if you do not know what you are doing, you risk getting completely lost.

When working with **sealert**, you can see that different plug-ins are called, and every plug-in has a confidence score. If, as in the example in Example 22-6, one plug-in has an 83.8% confidence score, while the other has only a 17.1% confidence score, it may be obvious that the former approach is what you should choose. Unfortunately, however, it is not always that readable.

TIP If you are not sure what SELinux is trying to tell you, install the **setroubleshoot-server** package and analyze what **sealert** shows. The information that is shown by **sealert** is often a lot more readable. Sometimes it will not help you at all, whereas sometimes the information can prove quite helpful.

Preparing for SELinux Coverage on the Exam

If this is the first time you’ve read about SELinux, you may find it overwhelming. That’s natural, because SELinux is rather overwhelming. In this chapter I’ve tried to give you a rather complete overview of how SELinux works. Fortunately, you don’t have to know all of this material for the RHCSA exam.

In the current specification of the RHCSA exam, there is nothing about configuration of services. SELinux is mostly useful for configuring services, but you don’t have to install and secure any web, FTP, or file services on the RHCSA 8 exam. What does remain are the following tasks, and you better make sure you master them well before taking the exam:

- Ensure that SELinux is enabled and in enforcing mode by editing `/etc/sysconfig/selinux`.
- Use **restorecon** to reapply the right context to a file or directory.
- Use **sealert** to troubleshoot why SELinux isn’t working.
- Make sure you know how to troubleshoot SELinux port-related issues.

Summary

This chapter provided an RHCSA-level introduction to SELinux. You’ve learned why SELinux is needed for security and how SELinux uses context as the main feature to apply security. You’ve also learned how to set the default SELinux mode and how to analyze in case things go wrong.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 22-3 lists a reference of these key topics and the page number on which each is found.

Table 22-3 Key Topics for Chapter 22

Key Topic Element	Description	Page
Table 22-2	SELinux core elements	475
List	Elements a context label can be applied to	479
List	Three parts of a context label	480
List	How new context settings are applied	486



Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

policy, enforcing, permissive, context, context type, source context, target context, audit log

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. You want to put SELinux temporarily in permissive mode. Which command do you use?
2. You need a list of all available Booleans. Which command do you use?
3. You do not see any service-specific SELinux man page. What solution do you need to apply?
4. What is the name of the package you need to install to get easy-to-read SELinux log messages in the audit log?
5. What commands do you need to run to apply the httpd_sys_content_t context type to the directory /web?

6. When would you use the **chcon** command?
7. Which file do you need to change if you want to completely disable SELinux?
8. Where does SELinux log all of its messages?
9. You have no clue which context types are available for the ftp service. What command enables you to get more specific information?
10. Your service does not work as expected, and you want to know whether it is due to SELinux or something else. What is the easiest way to find out?

End-of-Chapter Lab

You have now learned how SELinux works. To practice managing this essential service, work through this end-of-chapter lab about SELinux.

Lab 22.1

1. Change the Apache document root to **/web**. In this directory, create a file with the name **index.html** and give it the contents **welcome to my web server**. Restart the httpd process and try to access the web server. This will not work. Fix the problem.
2. In the home directory of the user root, create a file with the name **hosts** and give it the following contents:

```
192.168.4.200 labipa.example.com
192.168.4.210 server1.example.com
192.168.4.220 server2.example.com
```
3. Move the file to the /etc directory and do what is necessary to give this file the correct context.

This page intentionally left blank



The following topics are covered in this chapter:

- Understanding Linux Firewalling
- Working with Firewalld

The following RHCSA exam objective is covered in this chapter:

- Restrict network access using firewall-cmd/firewall

If a server is connected to the Internet, it needs to be protected against unauthorized access. SELinux is one part of this protection as discussed in Chapter 22, “Managing SELinux”), and a firewall is the second part. The Linux kernel implements firewalling via the netfilter framework. To configure which packets are allowed and which are not, Firewalld is the default solution in RHEL 8. In this chapter, you learn how a basic Firewalld configuration is created in an RHEL 8 environment.

Configuring a Firewall

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 23-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 23-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Understanding Linux Firewalling	1–3, 7
Working with Firewalld	4–6, 8–10

1. Which of the following is not a standard Firewalld zone?
 - a. untrusted
 - b. trusted
 - c. external
 - d. internal
2. Which of the following is the name of the firewalling service as implemented in the Linux kernel?
 - a. iptables
 - b. firewalld
 - c. netfilter
 - d. firewall-mod

3. Which of the following is *not* an advantage of Firewalld?
 - a. Rules can be modified through dbus.
 - b. It has an easy-to-use command-line interface.
 - c. It has an easy-to-use graphical interface.
 - d. It can be used to manage the iptables service.
4. Which command enables you to list all available Firewalld services?
 - a. **firewall-cmd --list-services**
 - b. **firewall-cmd --list-all**
 - c. **firewall-cmd --get-services**
 - d. **firewall-cmd --show-services**
5. What is the name of the GUI tool that enables you to easily manage Firewalld configurations?
 - a. **system-config-firewall**
 - b. **firewall-gtk**
 - c. **firewall-config**
 - d. **firewall-gui**
6. Which of the following shows the correct syntax for adding a port persistently to the current Firewalld configuration?
 - a. **firewall-cmd --addport=2022/tcp --permanent**
 - b. **firewall-cmd --add-port=2022/tcp --permanent**
 - c. **firewall-cmd --addport=2022/tcp --persistent**
 - d. **firewall-cmd --add port=2022/tcp --persistent**
7. Which zone should you use for an interface that is on a network where you need minimal firewall protection because every other computer on that same network is trusted?
 - a. Trusted
 - b. Home
 - c. Work
 - d. Private

8. Which of the following statements is true about the **--permanent** command-line option when used with **firewall-cmd**?
 - a. Configuration that is added using **--permanent** is activated immediately and will be activated automatically after (re)starting **Firewalld**.
 - b. Configuration that is added using **--permanent** is activated immediately.
 - c. Configuration that is added using **--permanent** is not activated immediately and can be activated only by using **systemctl restart firewalld**.
 - d. To activate configuration that has been added with the **--permanent** option, you need to reload the firewall configuration by using **firewall-cmd --reload**.
9. Which command enables you to get an overview of all the current firewall configurations for all zones?
 - a. **firewall-cmd --show-current**
 - b. **firewall-cmd --list-all**
 - c. **firewall-cmd --list-current**
 - d. **firewall-cmd --show-all**
10. Which of the following statements is *not* true about the **firewall-config** GUI tool?
 - a. All configuration that is created in **firewall-config** is automatically activated and stored permanently.
 - b. The **firewall-config** tool provides an easy-to-use interface to add ports to zones.
 - c. In its default screen, **firewall-config** shows all zones.
 - d. **firewall-config** connects to the **Firewalld** service. If this service is not running, you may have problems working with **firewall-config**.

Foundation Topics

Understanding Linux Firewalling

Key Topic

You can use a firewall to limit traffic coming in to a server or going out of the server. Firewalling is implemented in the Linux kernel by means of the netfilter subsystem. Netfilter allows kernel modules to inspect every incoming, outgoing, or forwarded packet and act upon such a packet by either allowing it or blocking it. So, the kernel firewall allows for inspection of incoming packets, outgoing packets, and packets that are traversing from one interface to another if the RHEL server is providing routing functionality.

Understanding Previous Solutions

To interact with netfilter, different solutions can be used. On earlier versions of Red Hat Enterprise Linux, iptables was the default solution to configure netfilter packet filtering. This solution worked with the command-line utility **iptables**, which provided a sophisticated and detailed way of defining firewall rules, but that also was challenging to use for the occasional administrator because of the complicated syntax of **iptables** commands and because the ordering rules could become relatively complex.

The iptables service is no longer offered in Red Hat Enterprise Linux 8. It has been replaced with **nftables**, a newer solution with more advanced options than the ones offered by iptables. The **nftables** command-line tool offers an advanced interface to write rules directly to nftables.

Understanding FirewallD

Key Topic

In Red Hat Enterprise Linux 7 a new firewall solution was introduced: FirewallD is a system service that can configure firewall rules by using different interfaces. Administrators can manage rules in a FirewallD environment, but even more important is that applications can request ports to be opened using the Dbus messaging system, which means that rules can be added or removed without any direct action required of the system administrator, which allows applications to address the firewall from user space.

FirewallD was developed as a completely new solution for managing Linux firewalls. It uses the FirewallD service to manage the netfilter firewall configuration.

Understanding Firewalld Zones

Firewalld makes firewall management easier by working with zones. A *zone* is a collection of rules that are applied to incoming packets matching a specific source address or network interface. Firewalld applies to incoming packets only by default, and no filtering happens on outgoing packets.

The use of zones is particularly important on servers that have multiple interfaces. On such servers, zones allow administrators to easily assign a specific set of rules. On servers that have just one network interface, you might very well do with just one zone, which is the default zone. Every packet that comes into a system is analyzed for its source address, and based on that source address, Firewalld analyzes whether or not the packet belongs to a specific zone. If not, the zone for the incoming network interface is used. If no specific zone is available, the packet is handled by the settings in the default zone.

Key
Topic

Firewalld works with some default zones. Table 23-2 describes these default zones.

Table 23-2 Firewalld Default Zones

Key
Topic

Zone Name	Default Settings
block	Incoming network connections are rejected with an “icmp-host-prohibited” message. Only network connections that were initiated on this system are allowed.
dmz	For use on computers in the demilitarized zone. Only selected incoming connections are accepted, and limited access to the internal network is allowed.
drop	Any incoming packets are dropped and there is no reply.
external	For use on external networks with masquerading (Network Address Translation [NAT]) enabled, used especially on routers. Only selected incoming connections are accepted.
home	For use with home networks. Most computers on the same network are trusted, and only selected incoming connections are accepted.
internal	For use in internal networks. Most computers on the same network are trusted, and only selected incoming connections are accepted.
public	For use in public areas. Other computers in the same network are not trusted, and limited connections are accepted. This is the default zone for all newly created network interfaces.
trusted	All network connections are accepted.
work	For use in work areas. Most computers on the same network are trusted, and only selected incoming connections are accepted.



Understanding Firewalld Services

The second key element while working with Firewalld is the service. Note that a service in Firewalld is *not* the same as a service in systemd; a Firewalld service specifies what exactly should be accepted as incoming and outgoing traffic in the firewall. It typically includes ports to be opened, as well as supporting kernel modules that should be loaded. In Firewalld, some default services are defined, which allows administrators to easily allow or deny access to specific ports on a server.

Behind each service is a configuration file that explains which UDP or TCP ports are involved and, if so required, which kernel modules must be loaded. To get a list of all services available on your computer, you can use the command **firewall-cmd --get-services** (see Example 23-1).

Example 23-1 Use **firewall-cmd --get-services** for a List of All Available Services

```
[root@server1 ~]# firewall-cmd --get-services
RH-Satellite-6 amanda-client amanda-k5-client amqp amqps apcupsd
audit bacula bacula-client bgp bitcoin bitcoin-rpc bitcoin-testnet
bitcoin-testnet-rpc ceph ceph-mon cfengine cockpit condor-collector
ctdb dhcp dhcpv6 dhcpv6-client distcc dns docker-registry docker-
swarm dropbox-lansync elasticsearch etcd-client etcd-server finger
freeipa-ldap freeipa-ldaps freeipa-replication freeipa-trust ftp
ganglia-client ganglia-master git gre high-availability http https
imap imaps ipp ipp-client ipsec irc ircs iscsi-target isns jenkins
kadmin kerberos kibana klogin kpasswd kprop kshell ldap ldaps
libvirt libvirt-tls lightning-network llmnr managesieve matrix mdns
minidlna mongodb mosh mounsd mqtt mqtt-tls ms-wbt mssql murmur
mysql nfs nfs3 nmea-0183 nrpe ntp nut openvpn ovirt-imageio ovirt-
storageconsole ovirt-vmconsole plex pmcd pmproxy pmwebapi pmwebapis
pop3 pop3s postgresql privoxy proxy-dhcp ptp pulseaudio puppetmaster
quassel radius redis rpc-bind rsh rsyncd rtsp salt-master samba
samba-client samba-dc sane sip sips slp smtp smtp-submission smtps
snmp snmptrap spideroak-lansync squid ssh steam-streaming svdrp svn
syncthing syncthing-gui synergy syslog syslog-tls telnet tftp tftp-
client tinc tor-socks transmission-client upnp-client vdsd vnc-
server wbem-http wbem-https wsman wsmans xdmcp xmpp-bosh xmpp-client
xmpp-local xmpp-server zabbix-agent zabbix-server
```

In essence, what it comes down to when working with Firewalld is that the right services need to be added to the right zones. In special cases, the configuration may be enhanced with more specific settings. In the next section, you learn which tools you can use for that purpose.

Behind each service is an XML configuration file. Changes are easily made to these XML files. Default (RPM installed) XML files are stored in `/usr/lib/firewalld/services`. Custom XML files can be added to the `/etc/firewalld/services` directory and will automatically be picked up after restarting the Firewalld service.

Example 23-2 shows what the contents of a service file looks like.

Example 23-2 Contents of the ftp.xml Service File

```
[root@server1 services]# cat ftp.xml
<?xml version="1.0" encoding="utf-8"?>
<service>
  <short>FTP</short>
  <description>FTP is a protocol used for remote file transfer.
  If you plan to make your FTP
  server publicly available, enable this option. You need the vsftpd
  package installed for this
  option to be useful.</description>
  <port protocol="tcp" port="21"/>
  <module name="nf_conntrack_ftp"/>
</service>
```

Working with Firewall

In this section, you learn how to configure a firewall with the Firewalld command-line interface tool, **firewall-cmd**. The Firewalld service also offers a GUI version of this tool, **firewall-config**, but the RHCSA exam objectives list only **firewall-cmd**, so this section focuses on working from the command line. This easily accessible tool enables uncomplicated configuration.

When working with either of these tools, be aware of where exactly modifications are made. Both tools work with an in-memory state of the configuration in addition to an on-disk state (permanent state) of the configuration. While using either of these tools, make sure to commit changes to disk before proceeding.

The **firewall-cmd** tool is an easily accessible tool that enables administrators to change the runtime configuration of the firewall and to write this configuration to disk. Before learning all the options available with this versatile command, in Exercise 23-1 you work with some of the most important options **firewall-cmd** offers.

Exercise 23-1 Managing the Firewall with **firewall-cmd**

1. Open a root shell. Type **firewall-cmd --get-default-zone**. This shows the current default zone. You'll see the current default zone, which is by default set to public.
2. To see which zones are available, type **firewall-cmd --get-zones**.

3. Show the services that are available on your server by typing **firewall-cmd --get-services**. Notice that the **firewall-cmd --get** options show what is available on your server, so basically you can use **firewall-cmd --get-<item>** to request information about a specific item.
4. To see which services are available in the current zone, type **firewall-cmd --list-services**. You'll see a short list containing a Dynamic Host Configuration Protocol (DHCP) client as well as Secure Shell (SSH) and the cockpit web-based management interface.
5. Type **firewall-cmd --list-all**. Look at the output and compare the output to the result of **firewall-cmd --list-all --zone=public**. Both commands show a complete overview of the current firewall configuration, as shown in Example 23-3. Notice that you see much more than just the zone and the services that are configured in that zone; you also see information about the interfaces and more advanced items.

Example 23-3 Showing Current Firewall Configuration

```
[root@server1 ~]# firewall-cmd --list-all
public (active)
  target: default
  icmp-block-inversion: no
  interfaces: ens160
  sources:
  services: cockpit dhcpv6-client ssh
  ports:
  protocols:
  masquerade: no
  forward-ports:
  source-ports:
  icmp-blocks:
  rich rules:
```

6. Type **firewall-cmd --add-service=vnc-server** to add the VNC server to the configuration of the firewall. Verify using **firewall-cmd --list-all**.
7. Type **systemctl restart firewalld** and repeat **firewall-cmd --list-all**. Notice that the vnc-server service is no longer listed; this is because the previous command has added the service to the runtime configuration but not to the persistent configuration.

8. Add the `vnc-server` service again, but make it permanent this time, using **`firewall-cmd --add-service vnc-server --permanent`**.
9. Type **`firewall-cmd --list-all`** again to verify. You'll see that VNC server service is not listed. Services that have been added to the on-disk configuration are not added automatically to the runtime configuration. Type **`firewall-cmd --reload`** to reload the on-disk configuration into the runtime configuration.
10. Type **`firewall-cmd --add-port=2022/tcp --permanent`**, followed by **`firewall-cmd --reload`**. Verify using **`firewall-cmd --list-all`**. You'll see that a port has now been added to the Firewall configuration.

TIP On the exam, work with services as much as possible. Only use specific ports if no services contain the ports that you want to open.

In the preceding exercise, you worked with zones and services and you learned how to add services and ports to the default zone. You should work with services as much as possible; adding individual ports is not recommended practice. The **`firewall-cmd`** interface offers many more options. Table 23-3 describes some of the most important command-line options.

Table 23-3 Common **`firewall-cmd`** Options

firewall-cmd Options	Explanation
<code>--get-zones</code>	Lists all available zones
<code>--get-default-zone</code>	Shows the zone currently set as the default zone
<code>--set-default-zone=<ZONE></code>	Changes the default zone
<code>--get-services</code>	Shows all available services
<code>--list-services</code>	Shows services currently in use
<code>--add-service=<service-name></code> <code>[--zone=<ZONE>]</code>	Adds a service to the current default zone or the zone that is specified
<code>--remove-service=<service-name></code>	Removes a service from the configuration
<code>--list-all [--zone=<ZONE>]</code>	Lists all configurations in a zone
<code>--add-port=<port/protocol></code> <code>[--zone=<ZONE>]</code>	Adds a port and protocol
<code>--remove-port=<port/protocol></code> <code>[--zone=<ZONE>]</code>	Removes a port from the configuration
<code>--add-interface=<INTERFACE></code> <code>[--zone=<ZONE>]</code>	Adds an interface to the default zone or a specific zone that is specified

**Key
Topic**

firewall-cmd Options	Explanation
<code>--remove-interface=<INTERFACE></code> <code>[--zone=<ZONE>]</code>	Removes an interface from a specific zone
<code>--add-source=<ipaddress/netmask></code> <code>[--zone=<ZONE>]</code>	Adds a specific IP address
<code>--remove-source=<ipaddress/netmask></code> <code>[--zone=<ZONE>]</code>	Removes an IP address from the configuration
<code>--permanent</code>	Writes configuration to disk and not to runtime
<code>--reload</code>	Reloads the on-disk configuration

Summary

In this chapter, you learned how to set up a basic firewall environment, where Firewalld services are added to Firewalld zones to allow access to specific services on your computer. You also learned how to set up a base firewall by using the **firewall-cmd** command-line tool.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 23-4 lists a reference of these key topics and the page number on which each is found.



Table 23-4 Key Topics for Chapter 23

Key Topic Element	Description	Page
Paragraph	Introduces firewalling in the Linux kernel	500
Paragraph	Introduces netfilter as opposed to other firewalling tools	500
Paragraph	Introduces how Firewalld zones are used	501
Table 23-2	Firewalld default zones	501
Section	Introduces the concept of a Firewalld service	502
Table 23-3	Common firewall-cmd options	505

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

firewall, netfilter, iptables, nftables, Firewallld, zones, services

Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. Which service should be running before you try to create a firewall configuration with **firewall-config**?
2. Which command adds UDP port 2345 to the firewall configuration in the default zone?
3. Which command enables you to list all firewall configuration in all zones?
4. Which command enables you to remove the vnc-server service from the current firewall configuration?
5. Which **firewall-cmd** command enables you to activate a new configuration that has been added with the **--permanent** option?
6. Which **firewall-cmd** option enables you to verify that a new configuration has been added to the current zone and is now active?
7. Which command enables you to add the interface eno1 to the public zone?
8. If you add a new interface to the firewall configuration while no zone is specified, which zone will it be added to?
9. Which command enables you to add the source IP address 192.168.0.0/24 to the default zone?
10. Which command enables you to list all services that are currently available in Firewallld?

End-of-Chapter Lab

You have now learned how to work with Firewalld on a Red Hat Enterprise Linux 8 server. Make sure to master these skills by working through this end-of-chapter lab.

Lab 23.1

1. Create a firewall configuration that allows access to the following services that may be running on your server:
 - web
 - ftp
 - ssh
2. Make sure the configuration is persistent and will be activated after a restart of your server.

This page intentionally left blank



The following topics are covered in this chapter:

- Using NFS Services
- Using CIFS Services
- Mounting Remote File Systems Through fstab
- Using Automount to Mount Remote File Systems

The following RHCSA exam objective is covered in this chapter:

- Mount and unmount network file systems using NFS

The RHCSA exam requires that you know how to access network storage. This encompasses different topics. We'll discuss accessing network storage that has been provided through CIFS and NFS. You'll learn how to mount network storage through the fstab file, as well as how to automatically mount this storage using automount.

Accessing Network Storage

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 24-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 24-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Using NFS Services	1–4
Using CIFS Services	5
Mounting Remote File Systems Through fstab	6
Using Automount to Mount Remote File Systems	7–10

1. If you want to enable an NFS share where all access to files in the share is anonymous, which of the following security options should you use?
 - a. You do not have to specify anything; the default setting is based on anonymous access to files.
 - b. none
 - c. sys
 - d. krb5
2. Which of the following is not a new feature in NFSv4?
 - a. Integration with Active Directory
 - b. Kerberized security
 - c. Services offered on TCP port 2049
 - d. The root mount

3. What is the name of the package that needs to be installed to mount NFS shares on an NFS client?
 - a. nfs-client
 - b. nfs-tools
 - c. nfs-utils
 - d. nfs
4. You type the command **showmount -e** to display available mounts on an NFS server, but you do not get any result. Which of the following is the most likely explanation?
 - a. The NFS client software is not running.
 - b. You are using a UID that does not exist on the server.
 - c. SELinux is not configured properly.
 - d. The firewall does not allow **showmount** traffic.
5. You want to log in to an SMB share. Which of the following commands shows correct syntax for doing so?
 - a. **mount -o username=sambauser1 //server/share /somewhere**
 - b. **mount -o uname=sambauser1 //server/share /somewhere**
 - c. **mount sambauser1@//server/share /somewhere**
 - d. **mount -o username=sambauser1@//server/share /somewhere**
6. Which of the following statements about authentication is true when configuring automatic mounts of SMB shares through fstab?
 - a. You need to specify the username; while the share is mounted, you'll see a prompt for a password.
 - b. You can only specify a username and a password by using mount options in /etc/fstab.
 - c. You do not have to specify a username or a password in fstab; you are prompted for them when accessing the share.
 - d. A secure way to specify a username and a password is by using a credentials file while mounting the share.

7. Which of the following is not a required step in configuring automount?
 - a. Identify the name of the automount directory in `/etc/auto.master`.
 - b. Create an indirect file in `/etc/auto.something`.
 - c. Start and enable the `autofs` service.
 - d. On the local mount point, set the appropriate permissions.
8. Assuming that the name of the directory you want automount to monitor is `/myfiles`, what is the recommended name for the corresponding configuration file?
 - a. `/etc/automount/auto.myfiles`
 - b. `/etc/auto.myfiles`
 - c. `/etc/myfiles.auto`
 - d. There is no recommended name.
9. Which of the following lines correctly identifies the syntax of a wildcard automount configuration that uses the NFS protocol?
 - a. `&. -rw server:/homes/*`
 - b. `&. rw. server:/homes/*`
 - c. `* -rw server:/homes/&`
 - d. `* rw. server:/homes/&`
10. What is the name of the service that automount uses?
 - a. `autofs`
 - b. `automount`
 - c. `autofiles`
 - d. `auto`

Foundation Topics

Using NFS Services

In previous chapters, you learned how to work with local file systems and mount them into the file system structure. In this chapter, you learn how to work with network file systems. The classic network file system is NFS (which stands for Network File System). It is a protocol that was developed for UNIX by Sun in the early 1980s, and it has been available on Linux forever. Its purpose is to make it possible to mount remote file systems into the local file system hierarchy.

Understanding NFS Security

When NFS was developed in the 1980s, it was often used together with Network Information Service (NIS), a solution that provides a network-based authentication server. With the use of NIS, all servers connected to NIS used the same user accounts and security was dealt with by the NIS server. The only thing that needed to be configured on the NFS server was host access. So, NFS security by default was limited to allowing and restricting specific hosts to access it.

Since the 1990s, NIS is not used often anymore. NFS, however, continues to be a very popular service, primarily because it is fast and easy to configure. Without NIS, the feature that provided user-based security has been removed, and that may make NFS seem to be an unsecure solution. Let's have a look at an example: Imagine that on server1, user linda has UID 1001. On server2, which is the NFS server, UID 1001 is used by user bob. After successfully connecting from server1 to server2, server1 user linda would have the same access to server2 resources as user bob. This obviously is an undesired situation.

To prevent situations like this from happening, NFS should be used together with a centralized authentication service. Commonly, a combination of the Lightweight Directory Access Protocol (LDAP) and Kerberos is used to provide this functionality. Configuration and integration of NFS with LDAP and Kerberos is not included in the RHCSA exam objectives, and for that reason will not be covered here.

RHEL 8 NFS Versions

On Red Hat Enterprise Linux 8, NFS 4 is the default version of NFS. If when making an NFS mount the NFS server offers a previous version of NFS, the client falls automatically back to that version. From a client, you can also force a specific

NFS version to be used for the mount, by using the **mount** option **nfsvers=**. This can prove useful if you are connecting to a server or a device that offers NFS 3 only. Fortunately, this type of server or device is increasingly uncommon nowadays.

Mounting the NFS Share

Setting up an NFS server is not a part of the RHCSA exam. However, to practice your NFS-based skills, it's useful to set up your own NFS test server. To do so, you need to go through a few tasks:

1. Create a local directory you want to share.
2. Edit the `/etc/exports` file to define the NFS share.
3. Start the NFS server.
4. Configure your firewall to allow incoming NFS traffic.

Exercise 24-1 guides you through these steps.

Exercise 24-1 Offering an NFS Share

You need a second server to do this exercise. A RHEL 8 server that was installed using the minimal server installation pattern is sufficient. This exercise assumes that a server with the name `server2.example.com` is available to offer these services.

1. Type **`mkdir -p /nfsdata /users/user1 /users/user2`** to create some local directories that are going to be shared.
2. Copy some random files to this directory, using **`cp /etc/[a-c]* /nfsdata`**.
3. Use **`vim`** to create the `/etc/exports` file and give it the following contents:

```
/nfsdata.    *(rw,no_root_squash)
/users.      *(rw,no_root_squash)
```
4. Type **`yum install -y nfs-utils`** to install the required packages.
5. Type **`systemctl enable --now nfs-server`** to start and enable the NFS server.
6. Type **`firewall-cmd --add-service nfs --permanent`** to add the NFS service. Also type **`firewall-cmd --add-service rpb-bind --permanent`** and **`firewall-cmd --add-service mountd --permanent`** to add the bind and mountd services.
7. To make the newly added services effective at this point, type **`firewall-cmd --reload`**.

Mounting the NFS Share

To mount an NFS share, you first need to find the names of the shares. This information can be provided by the administrator, but it is also possible to find out yourself. To discover which shares are available, you have multiple options:

- If NFSv4 is used on the server, you can use a root mount. That means that you just mount the root directory of the NFS server, and under the mount point you'll only see the shares that you have access to.
- Use the **showmount -e nfsserver** command to find out which shares are available.

WARNING The **showmount** command may have issues with NFSv4 servers that are behind a firewall. This is because **showmount** relies on the portmapper service, which uses random UDP ports while making a connection, and the firewalld nfs service opens port 2049 only, which does not allow portmapper traffic. If the firewall is set up correctly, the mountd and rpc-bind services need to be added to the firewall as well. It is very well possible that shares have been set up correctly on the server but you cannot see them because **showmount** does not get through the firewall. If you suspect that this is the case, use the NFS root mount as explained in Exercise 24-2, or just try mounting the NFS share.

Exercise 24-2 Mounting an NFS Share

1. On server1, type **yum install -y nfs-utils** to install the RPM package that contains the **showmount** utility.
2. Type **showmount -e server2.example.com** to see all exports available from server2.
3. On server1, type **mount server2.example.com:/mnt**. (Note the space between the slashes in the command.) This performs an NFSv4 pseudo root mount of all NFS shares.
4. Type **mount | grep server2** to verify the mount has succeeded.
5. Still on server1, type **ls /mnt**. This shows the subdirectories data and home, which correspond to the mounts offered by the NFS server.

Using CIFS Services

In the 1990s, Microsoft published the technical specifications of its Server Message Block (SMB) protocol. This protocol is the foundation of all shares that are created

in a Windows environment. Releasing these specifications led to the start of the Samba project. The goal of this project was to provide SMB services on top of other operating systems. Throughout the years, Samba has developed into the de facto standard for sharing files between different operating systems.

Samba is the standard Windows interoperability suite of programs for UNIX and Linux. As Samba has become such a common solution, it has been standardized and is now often referred to as the Common Internet File System (CIFS).

Discovering CIFS Shares

Before a CIFS share can be mounted, make sure that the `cifs-utils` and the `samba-client` RPM packages are installed on the client. After installing this, you can use the **`smbclient -L <servername>`** command to discover available Samba shares.

The **`smbclient`** command will ask for the password of the current user. This is because **`smbclient`** is a very generic utility that allows you to list shares but also to log in to Samba shares on remote servers and fetch files from the remote server. To list shares, however, no credentials are required. So, when the command asks for the password, just press Enter. Example 24-1 shows the result of the **`smbclient -L`** command.

Example 24-1 Discovering CIFS Shares

```
[root@server1 ~]# smbclient -L server2.example.com
Enter SAMBA\root's password:
Anonymous login successful
```

Sharename	Type	Comment
-----	----	-----
print\$	Disk	Printer Drivers
smbashare	Disk	smbashare
IPC\$	IPC	IPC Service (Samba 4.9.1)

```
Reconnecting with SMB1 for workgroup listing.
Anonymous login successful
```

Server	Comment
-----	-----
Workgroup	Master
-----	-----

In Example 24-1, you can see how the **smbclient** utility is used to discover available shares on IP address 192.168.4.201. (To bypass possible problems in name resolution, the IP address of the host is used instead of the name.) The command then prompts for the password of the current user. Next, the current domain or workgroup of the Samba server is shown, as well as the available shares.

Mounting Samba Shares

To mount a Samba share, you can use the **mount** command. You can use the **-t cifs** option to specify that the mount is to a Samba share, but without this option it will also work because the **mount** command is smart enough to discover by itself that it is a Samba share you want to connect to. If guest access is allowed on the share, you can specify the **-o guest** option to authenticate as the guest user without a password. Alternatively, use the **-o user=guest** option for the same purpose. The complete **mount** command looks like this:

```
mount -t cifs -o user=guest //192.168.4.200/data /mnt
```

This command mounts the /data share that is available on 192.168.4.200 on the local /mnt directory. Note that you'll be able to access files in the share but not write files in the share, because you are authenticated as the guest user, who has limited access permissions to the share.

Authenticating to Samba Shares

In the preceding section, you read how to mount a share with guest credentials. If you want to do something with the share, you should authenticate as a valid Samba user. This is a specific user account that has the credentials required to connect to a Samba share. (An ordinary Linux user cannot do that.) For this purpose, a Samba user must have been created.

To specify the Samba username you want to use, you can add the **-o username=someone** mount option:

```
mount -o username=sambauser1 //server/share /somewhere
```

When you do this, the **mount** command prompts for a password.

Configuring a Samba Server

For the RHCSA exam, you won't have to configure a Samba server. To allow you to practice Samba in your own environment, Exercise 24-3 guides you through

the steps that are required for setting up a Samba server. Perform these steps on a separate server with the name `server2.example.com`. Make sure to set up hostname resolution by using `/etc/hosts` so that this server can be reached by its name from `server1.example.com`.

Exercise 24-3 Setting Up a Samba Server

1. Log in to `server2` and open a root shell; mounts need to be performed as the root user.
2. Type **`mkdir /sambashare`** so that you have a directory that can be shared through Samba.
3. Open the share for SELinux using **`semanage fcontext -a -t public_content_t "/sambashare(/.*)?"`**.
4. Type **`restorecon -Rv /sambashare`** to apply the newly set SELinux context.
5. Copy some files to the share, using **`cp /etc/[fg]* /sambashare`**.
6. Enable the Linux-based access control by typing **`chmod 770 /sambashare`**.
7. Configure group ownership: **`chgrp sales /sambashare`**.
8. Install the Samba service by typing **`yum install -y samba`**.
9. Edit the `/etc/samba/smb.conf` configuration file and add the following lines:

```
[sambashare]
comment = sambashare
read only = No
path = /sambashare
```
10. Type **`systemctl enable --now smb`** to start and enable the Samba server.
11. Open the firewall by typing **`firewall-cmd --add-service samba`**.
12. To create a CIFS-compatible user account, type **`smbpasswd -a linda`**. This will add Samba credentials to the Linux user `linda`. Notice that this only works if you have an existing Linux user with the name `linda`.

In Exercise 24-4, you learn how to discover and mount Samba shares from the command line.

Exercise 24-4 Discovering and Mounting SMB Shares

This exercise assumes that you have set up a Samba server on `server2.example.com`. You must have completed Exercise 24-3 before you can do this exercise.

1. Log in to `server1` and open a root shell; mounts need to be performed as the root user.
2. Type `yum install -y cifs-utils samba-client` to install the required RPM packages.
3. Type `smbclient -L server2.example.com` to list available shares on the server.
4. Type `mount -t cifs -o username=linda //server2.example.com/sambashare /mnt` to mount the `/data` share as guest on the `/mnt` directory.
5. Type `mount` to verify that the mount has succeeded.

Mounting Remote File Systems Through `fstab`

You now know how to manually mount NFS and SMB file systems from the command line. If a file system needs to be available persistently, you need to use a different solution. Mounts can be automated either by using the `/etc/fstab` file or by using the `autofs` service. In this section, you learn how to make the mount through `/etc/fstab`. This is a convenient solution if you need the remote file system to be available permanently.

Mounting NFS Shares Through `fstab`

As you have learned in earlier chapters, the `/etc/fstab` file is used to mount file systems that need to be mounted automatically when a server restarts. Only the user root can add mounts to this configuration file, thus providing shares that will be available for all users. The `/etc/fstab` file can be used to mount the NFS file system as well as Samba. To mount an NFS file system through `/etc/fstab`, make sure that the following line is included:

```
server1:/share /nfs/mount/point nfs sync 0 0
```

When making an NFS mount through `fstab`, you have a few options to consider:



- In the first column, you need to specify the server and share name. Use a colon after the name of the server to identify the mount as an NFS share.
- The second column has the file system where you want to do the mount; this is not different from a regular mount.

- The third column contains the NFS file system type.
- The fourth column is used to specify mount options and includes the `sync` option. This ensures that modified files are committed to the remote file system immediately and are not placed in write buffers first (which would increase the risk of data getting lost).
- The fifth column contains a zero, which means that no backup support through the **dump** utility is requested.
- The sixth column also contains a zero, to indicate that no **fsck** has to be performed on this file system while booting to check the integrity of the file system. The integrity of the file system would need to be checked on the server, not on the client.

Mounting Samba Shares Through `fstab`

When mounting Samba file systems through `/etc/fstab`, you need to consider a specific challenge: You need to specify the user credentials that are needed to issue the mount. On an NFS share, this is not necessary because the user who accesses the shared file system by default does so using his own credentials. While mounting a Samba share through `/etc/fstab`, these user credentials are normally specified with the **username=** and **password=** mount options, but it is not a good idea to put these in clear text in the `/etc/fstab` file. The following line shows how you would automatically mount the file system that you created in Exercise 24-3:

```
//server2/sambashare    /sambamount.    cifs\    username=linda,  
password=password.    0 0
```

Using Automount to Mount Remote File Systems

As an alternative to using `/etc/fstab`, you can configure automount to mount the share automatically. Automount can be used for SMB as well as NFS mounts, and the big difference is that mounts through automount are affected on demand and not by default. So, using automount ensures that no file systems are mounted that are not really needed.

Understanding Automount

Automount is implemented by the `autofs` service that takes care of mounting a share when an attempt is made to access it. That means it is mounted on demand and that it does not have to be mounted permanently. An important benefit of using automount is that it works completely in user space and, contrary to mounts that are made through the **mount** command, no root permissions are required.

Defining Mounts in Automount

In automount, mounts are defined through a two-step procedure. First, you need to edit the master configuration file `/etc/auto.master`. In this directory you identify the mount point (for instance, `/nfsdata`). Next, and on the same line, you identify the name of the secondary file, as all further configuration happens in this secondary file. The line you create could look as follows:

```
/nfsdata.      /etc/auto.nfsdata
```

In the secondary file you put the name of the subdirectory that will be created in the mount point directory as a relative filename. For instance, you start the line with **files**, to mount `/nfsdata/files`. After the name of the subdirectory, you specify NFS mount options, as well as the server and share name to access the NFS share. This line could look as follows:

```
files.      -rw.      server2:/nfsdata
```

Configuring Automount for NFS

Configuring an automount solution is a multistep procedure. To show how it works, Exercise 24-5 lists all steps involved. Follow the steps in this exercise to see for yourself how to configure automount.

Exercise 24-5 Configuring Direct and Indirect Maps to Mount NFS Shares

This exercise is performed on `server1`. It uses the NFS shares provided by `server2` that you created in Exercise 24-1.

1. Type **yum install -y autofs** to install the autofs package.
2. Type **showmount -e server2.example.com**, which shows you NFS exports offered by `server2`.
3. Open the file `/etc/auto.master` and add the following line:
`/nfsdata/etc/auto.nfsdata`
4. Type **vim /etc/auto.nfsdata** and add the following line:
`files -rw server2:/nfsdata`
5. Type **systemctl enable --now autofs** to start and enable the autofs service.
6. Type **ls /**; notice that there is no `/nfsdata` directory.
7. Type **cd /nfsdata/files** to get access to the `/nfsdata` directory.
8. Type **mount** and notice the last three lines in the mount output, created by the autofs service.

Using Wildcards in Automount

In Exercise 24-5, you learned how to perform automounts based on fixed directory names. In some cases, this is not very useful, and you are better off using dynamic directory names. This is, for example, the case for automounting home directories.

With home directories, a very helpful solution is to have the home directory of a specific user automounted when that user logs in. So, for example, if user linda logs in, she gets access to the NFS exported directory /home/linda, and when user anna logs in, she gets access to /home/anna. Using wildcards in automount offers an excellent tool to do this.

To create a wildcard mount, you will use lines like `* -rw server2:/users/&`. In this line, the `*` represents the local mount point, which in this case represents anything, and the `&` represents the matching item on the remote server.

Obviously, you could also choose to export the /home directory and mount just the /home directory, but that increases the risk that user anna gets access to user linda's home directory. For that reason, using a wildcard mount is a much cleaner solution, as demonstrated in Exercise 24-6.

Exercise 24-6 Configuring Wildcard Mounts

This exercise is performed on server1. It uses the NFS shares that are provided by server2, which you created in Exercise 24-1. On server2, the directory /users is exported, which simulates an NFS server that exports home directories. You are going to configure a wildcard mount, such that when /users/user1 is accessed, that exact directory is mounted, and when /users/user2 is accessed, that directory is mounted.

1. Type **yum install -y autofs** to install the autofs package.
2. Open the file /etc/auto.master and make sure it includes the following line:
`/users. /etc/auto.users`
3. Create the file /etc/auto.users and give it the following contents:
`*. -rw. server2:/users/&`
4. Type **systemctl restart autofs** to restart the autofs service.
5. Type **cd /users/user1** to get access to the NFS export /users/user1 on the server2 server.

Summary

In this chapter, you learned how to mount remote file systems and how to set up an FTP server. You first learned how to manually mount either an NFS or an SMB file system from the command line. Then you learned how these mounts can be automated through `/etc/fstab` or `automount`. In the last section in this chapter, you learned how to set up an FTP server to allow for anonymous file uploads.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 24-2 lists a reference of these key topics and the page number on which each is found.



Table 24-2 Key Topics for Chapter 24

Key Topic Element	Description	Page
List	Options to consider when making an NFS mount through <code>fstab</code>	520

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

SMB, CIFS, NFS, `automount`

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. On your NFS server, you have verified that the `nfs` service is active, and the firewall allows access to TCP port 2049. A client uses **showmount** against your server but doesn't see any exports. What is the most likely explanation?
2. Which command enables you to show available NFS mounts on `server1`?
3. Which command enables you to mount an NFS share that is available on `server1:/share`?
4. Which command can you use to discover SMB mounts on a specific server?
5. Which package must be installed on an SMB client before you can make an SMB mount?
6. How do you mount the Samba share `data` on `server1` with guest access on the local directory `/mnt`?
7. How do you mount a Samba mount through `fstab` while avoiding putting the username and password in `/etc/fstab`?
8. What is the name of the main automount configuration file?
9. What is the name of the service that implements automount?
10. Which ports do you need to open in the firewall of the automount client?

End-of-Chapter Lab

In this chapter, you learned how to mount remote file systems and automate those mounts using `/etc/fstab` or `automount`. You also learned how to set up an FTP server and saw the essential parts of a typical FTP server. In this end-of-chapter lab, you practice these skills in a way that is similar to how you need to perform them on the exam.

Lab 24.1

1. Set up an NFS server that shares the `/home` directory on `server2`.
2. Configure `server1` to access the NFS-shared home directory using `automount`. You need to do this using wildcard `automount`.



The following topics are covered in this chapter:

- Understanding Local Time
- Using Network Time Protocol
- Managing Time on Red Hat Enterprise Linux

The following RHCSA exam objective is covered in this chapter:

- Configure time service clients

An increasing number of services offered through Linux servers depend on the correct configuration of time on the server. Think of services such as database synchronization, Kerberos authentication, and more. In this chapter, you learn how time is configured on a Linux server.

Configuring Time Services

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 25-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and ‘Review Questions.’”

Table 25-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Understanding Local Time	1–2
Using Network Time Protocol	4–5
Managing Time on Red Hat Enterprise Linux	3, 6–10

1. When a system is started, where does it initially get the system time?
 - a. NTP
 - b. Software time
 - c. The hardware clock
 - d. Network time
2. Which of the following statements is *not* true about local time?
 - a. Local time is the current time in the current time zone.
 - b. In local time, DST is considered.
 - c. System time typically should correspond to the current local time.
 - d. Hardware time typically corresponds to the current local time.

3. Which is the recommended command in RHEL 8 to set the local time zone?
 - a. **hwclock**
 - b. **tz**
 - c. **date**
 - d. **timedatectl**
4. Which clock type would you recommend on a server that is *not* connected to any other server but needs to be configured with the most accurate time possible?
 - a. RTC
 - b. UTC
 - c. An atomic clock
 - d. NTP
5. Which configuration file contains the default list of NTP servers that should be contacted on RHEL 8?
 - a. **/etc/ntp/ntp.conf**
 - b. **/etc/ntp.conf**
 - c. **/etc/chrony/chronyd.conf**
 - d. **/etc/chrony.conf**
6. Which of the following shows correct syntax to set the current system time to 9:30 p.m.?
 - a. **date 9:30**
 - b. **date --set 9.30 PM**
 - c. **date -s 21:30**
 - d. **date 2130**
7. Which command correctly translates epoch time into human time?
 - a. **time --date '@1420987251'**
 - b. **time --date '\$1420987251'**
 - c. **time --date '#1420987251'**
 - d. **time --date '1420987251'**

8. Which command enables you to monitor the difference between the hardware clock and system clock?
 - a. `tail -f /var/lib/time/drift`
 - b. `date -h`
 - c. `hwclock -c`
 - d. `hwclock -d`
9. Which command enables you to show current information that includes the local time, hardware time, and the time zone the system is in?
 - a. `timedatectl --all`
 - b. `timedatectl --tz`
 - c. `timedatectl -ht`
 - d. `timedatectl`
10. Which command can you use to verify that a time client that is running the chrony service has successfully synchronized?
 - a. `timedatectl`
 - b. `chronyc sources`
 - c. `systemctl chrony status`
 - d. `chronyc status`

Foundation Topics

Understanding Local Time

Key Topic

When a Linux server boots, the hardware clock, also referred to as the real-time clock, is read. This clock typically resides in the computer hardware, and the time it defines is known as hardware time. Generally, it is an integrated circuit on the system board that is completely independent of the current state of the operating system and keeps running even when the computer is shut down, as long as the mainboard battery or power supply feeds it. From the hardware clock, the system gets its initial time setting.

The time on the hardware clock on Linux servers is usually set to Coordinated Universal Time (UTC). UTC is a time that is the same everywhere on the planet, and based on UTC, the current local time is calculated. (Later in this chapter you learn how this works.)

Key Topic

System time is maintained by the operating system. Once the system has booted, the system clock is completely independent of the hardware clock. Therefore, when system time is changed, the new system time is not automatically synchronized with the hardware clock.

System time maintained by the operating system is kept in UTC. Applications running on the server convert system time into local time. Local time is the actual time in the current time zone. In local time, daylight saving time (DST) is considered so that it always shows an accurate time for that system. Table 25-2 gives an overview of the different concepts that play a role in Linux time.

Key Topic

Table 25-2 Understanding Linux Time

Concept	Explanation
Hardware clock	The hardware clock that resides on the main card of a computer system
Real-time clock	Same as the hardware clock
System time	The time that is maintained by the operating system
Software clock	Similar to system time
Coordinated Universal Time (UTC)	A worldwide standard time
Daylight saving time	Calculation that is made to change time automatically when DST changes occur
Local time	The time that corresponds to the time in the current time zone

Using Network Time Protocol

As you learned, the current system time is based on a hardware clock. This hardware clock is typically a part of the computer's motherboard, and it might be unreliable. Because of its potential unreliability, it is a good idea to use time from a more reliable source. Generally speaking, two solutions are available.

One option is to buy a more reliable hardware clock. This may be, for instance, a very accurate atomic clock connected directly to your computer. When such a very reliable clock is used, an increased accuracy of the system time is guaranteed. Using an external hardware clock is a common solution to guarantee that datacenter time is maintained, even if the connection to external networks for time synchronization temporarily is not available.

Another and more common solution is to configure your server to use Network Time Protocol (NTP). NTP is a method of maintaining system time that is provided through NTP servers on the Internet. It is an easy solution to provide an accurate time to servers, because most servers are connected to the Internet anyway.

To determine which Internet NTP server should be used, the concept of stratum is used. The stratum defines the reliability of an NTP time source, and the lower the stratum, the more reliable it is. Typically, Internet time servers are using stratum 1 or 2. When configuring local time servers, you can use a higher stratum number to configure the local time server as a backup, only it will never be used when Internet time is available.

It is good practice, for example, to set stratum 5 on a local time server with a very reliable hardware clock and stratum 8 on a local time server that is not very reliable. A setting of stratum 10 can be used for the local clock on every node that uses NTP time. This enables the server to still have synchronized time when no external connection is available. Stratum 15 is used by clocks that want to indicate they should not be used for time synchronization.

Setting up a server to use NTP time on RHEL 8 is easy if the server is already connected to the Internet. If this is the case, the `/etc/chrony.conf` file is configured with a standard list of NTP servers on the Internet that should be contacted. The only thing the administrator has to do is switch on NTP, by using **`timedatectl set-ntp 1`**.

**Key
Topic**

Managing Time on Red Hat Enterprise Linux

Different commands are involved in managing time on Red Hat Enterprise Linux. Table 25-3 provides an overview.

**Key
Topic****Table 25-3** Commands Related to RHEL 8 Time Management

Command	Short Description
date	Manages local time
hwclock	Manages hardware time
timedatectl	Developed to manage all aspects of time on RHEL 8

**Key
Topic**

On a Linux system, time is calculated as an offset of epoch time. *Epoch time* is the number of seconds since January 1, 1970, in UTC. In some logs (such as `/var/log/audit/audit.log`), you'll find time stamps in epoch time and not in human time. To convert such an epoch time stamp to human time, you can use the **--date** option, followed by the epoch string, starting with an **@**:

```
date --date '@1420987251'
```

The use of epoch time also creates a potential timing problem on Linux. On a 32-bit system, the number of seconds that can be counted in the field that is reserved for time notation will be exceeded in 2037. (Try setting the time to somewhere in 2050 if you are on a 32-bit kernel; it will not work.) 64-bit systems can address time until far into the twenty-second century.

Using date

The **date** command enables you to manage the system time. You can also use it to show the current time in different formats. Some common usage examples of **date** are listed here:

- **date**: Shows the current system time
- **date +%d-%m-%y**: Shows the current system day of month, month, and year
- **date -s 16:03**: Sets the current time to 3 minutes past 4 p.m.

Using hwclock

The **date** command enables you to set and show the current system time. Using the **date** command will not change the hardware time that is used on your system. To manage hardware time, you can use the **hwclock** command. The **hwclock** command has many options, some of which are of particular interest:

- **hwclock --systemd**: Synchronizes current system time to the hardware clock
- **hwclock --hctosys**: Synchronizes current hardware time to the system clock

Using **timedatectl**

A new command that was introduced in RHEL 7 that enables you to manage many aspects of time is **timedatectl**. As shown in Example 25-1, when used without any arguments, this command shows detailed information about the current time and date. It also displays the time zone your system is in, in addition to information about the use of NTP network time and information about the use of DST.

Example 25-1 Using **timedatectl** to Get Detailed Information About Current Time Settings

```
[root@server1 ~]# timedatectl
          Local time: Mon 2019-06-10 08:27:57 EDT
          Universal time: Mon 2019-06-10 12:27:57 UTC
          RTC time: Mon 2019-06-10 12:27:57
          Time zone: America/New_York (EDT, -0400)
System clock synchronized: yes
          NTP service: active
          RTC in local TZ: no
```

The **timedatectl** command works with commands to perform time operations. Table 25-4 provides an overview of the relevant commands.

Table 25-4 **timedatectl** Command Overview

Command	Explanation
status	Shows the current time settings
set-time TIME	Sets the current time
set-timezone ZONE	Sets the current time zone
list-timezone	Shows a list of all time zones
set-local-rtc [0 1]	Controls whether the RTC (the real-time clock, normally referred to as the hardware clock) is in local time
set-ntp [0 1]	Controls whether NTP is enabled

**Key
Topic**

The **timedatectl** command was developed as a generic solution to manage time on RHEL 7. It has some functions that are offered through other commands, but the purpose of the command is that eventually it will replace other commands used for managing time and date settings. When **timedatectl** is used to switch on NTP time, it talks to the **chronyd** process. Exercise 25-1 walks you through some common options to manage time on a RHEL 8 server.

Exercise 25-1 Managing Local Time

1. Open a root shell and type **date**.
2. Type **hwclock** and see whether it shows approximately the same time as **date** in step 1.
3. Type **timedatectl status** to show current time settings.
4. Type **timedatectl list-timezones** to show a list of all time zone definitions.
5. Type **timedatectl set-timezone Europe/Amsterdam** to set the current time zone to Amsterdam.
6. Type **timedatectl show** and note the differences with the previous output.
7. Type **timedatectl set-ntp 1** to switch on NTP use. You might see the error “failed to issue method call.” If you get this message, type **yum -y install chrony** and try again.
8. Open the configuration file **/etc/chrony.conf** and look up the server lines. These are used to specify the servers that should be used for NTP time synchronization.
9. Type **systemctl status chronyd** and verify that the chrony service is started and enabled. If this is not the case, use **systemctl start chronyd**; **systemctl enable chronyd** to make sure that it is operational.
10. Type **systemctl status -l chronyd** and read the status information.
Example 25-2 shows you what the output of the command should look like.

Example 25-2 Monitoring Current Time Synchronization Status

```
[root@server1 ~]# systemctl status -l chronyd
chronyd.service - NTP client/server
Loaded: loaded (/usr/lib/systemd/system/chronyd.service; enabled;
       vendor preset: enabled)
Active: active (running) since Mon 2019-06-10 05:22:30 EDT;
       3h 8min ago
Docs: man:chronyd(8)
      man:chrony.conf(5)
Main PID: 1062 (chronyd)
Tasks: 1 (limit: 11365)
Memory: 1.5M
CGroup: /system.slice/chronyd.service
        └─1062 /usr/sbin/chronyd
```

```
Jun 10 07:21:04 server1.example.com chronyd[1062]: Selected source
5.200.6.34
Jun 10 07:28:40 server1.example.com chronyd[1062]: Selected source
213.154.236.182
Jun 10 07:28:42 server1.example.com chronyd[1062]: Source
149.210.142.45 replaced with 195.242.98.57
Jun 10 07:43:51 server1.example.com chronyd[1062]: Selected source
5.200.6.34
Jun 10 07:53:35 server1.example.com chronyd[1062]: Selected source
195.242.98.57
Jun 10 08:16:24 server1.example.com chronyd[1062]: Forward time jump
detected!
Jun 10 08:16:24 server1.example.com chronyd[1062]: Can't synchronise:
no selectable sources
Jun 10 08:20:44 server1.example.com chronyd[1062]: Selected source
213.154.236.182
Jun 10 08:22:57 server1.example.com chronyd[1062]: Source
195.242.98.57 replaced with 195.191.113.251
Jun 10 08:25:05 server1.example.com chronyd[1062]: Selected source
5.200.6.34
```

Managing Time Zone Settings

Between Linux servers, time is normally communicated in UTC. This allows servers across different time zones to use the same time settings, which makes managing time in large organizations a lot easier. To make it easier for end users, though, the local time must also be set. To do this, the appropriate time zone needs to be selected.

On Red Hat Enterprise Linux 8, you have three approaches to setting the correct local time zone:

- Go to the directory `/usr/share/zoneinfo`, where you'll find different subdirectories containing files for each of the time zones that has been defined. To set the local time zone on a server, you can create a symbolic link with the name `/etc/localtime` to the time zone file that is involved. If you want to set local time to Los Angeles time, for instance, use **`ln -sf /usr/share/zoneinfo/America/Los_Angeles /etc/localtime`**.
- Use the **`tzselect`** utility. This tool starts the interface shown in Example 25-3, from which the appropriate region and locale can be selected.
- Use **`timedatectl`** to set the time zone information.

Example 25-3 Selecting the Time Zone Using **tzselect**

```
[root@localhost ~]# tzselect
Please identify a location so that time zone rules can be set
correctly.
Please select a continent, ocean, "coord", or "TZ".
 1) Africa
 2) Americas
 3) Antarctica
 4) Asia
 5) Atlantic Ocean
 6) Australia
 7) Europe
 8) Indian Ocean
 9) Pacific Ocean
10) coord - I want to use geographical coordinates.
11) TZ - I want to specify the time zone using the Posix TZ format.
#? 1
Please select a country whose clocks agree with yours.
 1) Algeria.           20) Gambia             39) Sao Tome & Principe
 2) Angola.            21) Ghana              40) Senegal
 3) Benin              22) Guinea            41) Sierra Leone
 4) Botswana           23) Guinea-Bissau     42) Somalia
 5) Burkina Faso       24) Kenya           43) South Africa
 6) Burundi           25) Lesotho           44) South Sudan
 7) C?te d'Ivoire      26) Liberia.         45) Spain
 8) Cameroon           27) Libya             46) St Helena
 9) Central African Rep. 28) Madagascar       47) Sudan
10) Chad               29) Malawi.           48) Swaziland
11) Comoros            30) Mali.             49) Tanzania
12) Congo (Dem. Rep.)  31) Mauritania.       50) Togo
13) Congo (Rep.)       32) Mayotte.          51) Tunisia
14) Djibouti           33) Morocco.         52) Uganda
15) Egypt              34) Mozambique.       53) Western Sahara
16) Equatorial Guinea  35) Namibia           54) Zambia
17) Eritrea            36) Niger.            55) Zimbabwe
18) Ethiopia           37) Nigeria
19) Gabon              38) Rwanda
#? 54
```

The following information has been given:

Zambia
Central Africa Time

Therefore TZ='Africa/Maputo' will be used.

Selected time is now: Mon Jul 22 12:03:41 CAT 2019.

Universal Time is now: Mon Jul 22 10:03:41 UTC 2019.

Is the above information OK?

1) Yes

2) No

#? 1

You can make this change permanent for yourself by appending the line
TZ='Africa/Maputo'; export TZ to the file '.profile' in your home
directory; then log out and log in again.

Here is that TZ value again, this time on standard output so that
you can use the /usr/bin/tzselect command in shell scripts:
Africa/Maputo

Configuring Time Service Clients

By default, the **chrony** service is configured to get the right time from the Internet. As a default configuration, the highly reliable time servers from pool.ntp.org are used to synchronize time. However, in a corporate environment it is not always desirable for time clients to go out to the Internet, and local time services should be used instead. This can be configured by making a simple modification to the chrony.conf configuration file.

By default, the chrony.conf configuration file contains the line pool 2.rhel.pool.ntp.org. If you comment out this line by putting a pound sign in front of it and add the line server yourtimeserver.example.com, your time server will be used instead of the servers in pool.ntp.org. Exercise 25-2 explains how to make this modification. Notice that this exercise requires access to two servers, in which server1 is configured as the time server and server2 is configured as the time client.

Exercise 25-2 Configuring an NTP Time Client

1. On server1, open a root shell.
2. Disable the line `pool 2.rhel.pool.ntp.org` by putting a `#` sign in front of it.
3. Include the line **`allow 192.168.0.0/16`** to allow access from all clients that use a local IP address starting with 192.168.
4. Also include the line **`stratum 8`**. This ensures that the local time server is going to advertise itself with a stratum of 8, which means it will be used by clients, but only if no Internet time servers are available. Next, close the configuration file.
5. Use **`systemctl restart chronyd`** to restart the chrony process with the new settings.
6. Still on server1, type **`firewall-cmd --add-service ntp --permanent`**, followed by **`firewall-cmd reload`**. This opens the firewall for time services.
7. Open a root shell on server2.
8. On server2, open the configuration file `/etc/chrony.conf` and disable the line **`pool 2.rhel.pool.ntp.org`**.
9. Add the line `server server1.example.com`. Make sure that name resolution to server1.example.com is configured, and if not, use the IP address of server1 instead.
10. Type **`systemctl restart chronyd`** to restart the chrony service with the new settings.
11. On server2, type the command **`chronyc sources`**. It should show the name or IP address of server1, the stratum of 8 that is advertised, and a synchronization status indicating that server2 has successfully synchronized its time.

Summary

In this chapter, you learned how time works on Linux. You read how your operating system can get its time by using hardware time, system time, and local time. You also learned how to manage time using the **`date`**, **`hwclock`**, and **`timedatectl`** commands.

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have several choices for exam preparation: the end-of-chapter labs; the memory tables in Appendix B; Chapter 26, “Final Preparation”; and the practice exams.

Review All Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 25-5 lists a reference of these key topics and the page number on which each is found.

Table 25-5 Key Topics for Chapter 25

Key Topic Element	Description	Page
Paragraph	Definition of hardware time	530
Paragraph	Definition of system time	530
Table 25-2	Understanding Linux time	530
Paragraph	Using NTP time	531
Table 25-3	Commands related to RHEL 8 time management	532
Paragraph	Explanation of epoch time	532
Table 25-4	timedatectl command overview	533

**Key
Topic**

Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key,” includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

hardware time, RTC, system time, network time, UTC, epoch time, time synchronization, stratum

Review Questions

The questions that follow are meant to help you test your knowledge of concepts and terminology and the breadth of your knowledge. You can find the answers to these questions in Appendix A.

1. Which command enables you to set the system time to 4:24 p.m.?
2. Which command sets hardware time to the current system time?
3. Which command enables you to show epoch time as human-readable time?
4. Which command enables you to synchronize the system clock with hardware time?
5. Which service is used to manage NTP time on RHEL 8?
6. Which command enables you to use NTP time on your server?
7. Which configuration file contains the list of NTP servers to be used?
8. Which command enables you to list time zones?
9. Which command enables you to set the current time zone?
10. How do you use chrony to set system time?

End-of-Chapter Lab

In this chapter, you learned how to manage time on Linux servers. Because it is very important to ensure that a server uses the correct time, you can now practice some of the most essential skills you have acquired in this chapter.

Lab 25.1

1. Compare the current hardware time to the system time. If there is a difference, make sure to synchronize time.
2. Set the time zone to correspond to the current time in Boston (USA East Coast).

Final Preparation

Congratulations! You made it through the book, and now it's time to finish getting ready for the RHCSA exam. This chapter helps you get ready to take and pass the exam. In this chapter, you learn more about the exam process and how to register for the exam. You also get some useful tips that will help you avoid some common pitfalls while taking the exam.

General Tips

In this section, you get some general tips about the exam. You learn how to verify your exam readiness, how to register for the exam, and what to do on the exam.

Verifying Your Readiness

Only register for the exam when you think that you are ready to pass it. This book contains a lot of material to help you verify your exam readiness. To start with, you should be able to answer all the “Do I Know This Already?” quiz questions, which you find at the beginning of each chapter. You should also have completed all the exercises in the chapters successfully, as well as the end-of-chapter labs. The labs are the first real way of testing your readiness because the questions are formulated similarly to the real exam questions, providing a good way to gauge whether you are ready for the exam.

Registering for the Exam

There are two ways of taking the RHCSA exam. You can take it either as a classroom exam or as a kiosk exam. The classroom exam is typically on Friday only, and it is offered primarily by Red Hat to provide an exam at the end of a course. Therefore, most of the people who are with you in the exam classroom have taken 4 days of course training before taking the exam, and for that reason, classroom exam availability is limited.

For a long time, the classroom exam was the only way to take the exam. For some time now, Red Hat has provided kiosk exams, too. Unlike a classroom exam, a kiosk exam is administered on an individual basis, where you work through the exam tasks on a kiosk computer. This is a monitored computer that is in a booth in an exam center, where you are monitored through multiple cameras while working on the exam tasks. The good thing about a kiosk exam is that you schedule the exam for a time and place that is convenient for your schedule.

You can register to take the exam either through Redhat.com (following the links under the Services & support tab) or through a training company. Either way, you take the same exam. It might be easier, though, to get a discount by booking through a local training company. Booking through Red Hat will be faster normally, as you have direct access to Red Hat.

If you book a classroom exam, you get an invitation for the time and date the exam is scheduled. If you book a kiosk exam, you get a voucher code that you can use to choose the exam venue, time, and date.

On Exam Day

Make sure to bring appropriate identification to the exam. To be allowed to take the exam, you need an approved government ID. Normally, a passport or driver's license will do; other types of proof may be accepted in your country as well. Do not forget it; without ID, you will not be allowed to take the exam.

Also, make sure you are on time. It is a good idea to arrive 30 minutes before the exam's scheduled starting time. If you are late, you will normally be allowed to sit for the exam, but you will not get extra time. So, just make sure that you are on time.

After proving your identity, you are introduced to the exam environment. Because of the nondisclosure agreement that every test-taker signs with Red Hat, I cannot tell you in detail what the exam environment looks like. I can tell you, though, that there will be an environment that consists of one or more servers. There is also a list of tasks that have to be performed. Work your way through the tasks, reading all carefully and thoroughly, and you will pass the exam if you have prepared well.

During the Exam

The tasks that you have to work on during the exam are not necessarily presented in the most logical order. Therefore, it is a good idea to start reading through all the tasks before you start working on the first assignment. While reading through all the tasks, you can decide which is the best order to create the configurations needed. Determine the best possible order for yourself, because it may not be obvious.

You have 2 hours and 30 minutes to work through all the exam assignments. Expect about 17 assignments; you might see a bit more or less.

Another very important tip is to read carefully, a skill that not many people have been taught well. IT administrators are very skilled in scanning web pages to retrieve the information that they need. That skill will not help you on the exam. Reading skills do. I cannot stress that enough. I estimate that 40% of all people who fail

the exam do so because they do not read the exam questions carefully. (They scan instead.) So, let me give you some tips on how to read the exam questions:

- If English is not your native language, you can switch the language that questions are presented in. Maybe the English-language question is not clear to you, but the question translated in another language is. So, if in doubt, read the translation as well.
- Because the questions are originally written in English—the default language and the preference of most exam candidates—they tend to be perfect in that form, because Red Hat applies a tremendous effort to make them perfect. Red Hat must rely on translators to ensure the questions are translated correctly, so the quality of the English-language questions is the best. You are free to use translated questions, but you should use the English-language questions as your primary source.
- To make sure that you do not miss anything, make a task list for each question. You have scratch paper with you during the exam. Use it to make a short list of tasks that you have to accomplish, and work on them one by one. This approach helps you concentrate on what the exam question is actually asking.
- After you have worked on all assignments, go have a cup of coffee. (You are allowed to take a break during the exam.) When you return, read all questions again to make sure that you did not miss anything. Taking a small break is important; it allows you to distance yourself from the exam, after which you should read the questions as if it is the first time that you have seen them.

Another important part of the exam is the order in which you work on the assignments. Regardless of the specific exam content, some topics need to be addressed before other topics. Make sure that you deal with those topics first. If you do not, it will make it more difficult or impossible to complete the other assignments. Roughly speaking, here is the order in which you should work on the exam topics:

1. Make sure that your server boots and you have root access to it.
2. Configure networking in the way it is supposed to work.
3. Configure any repositories that you need.
4. Install and enable all services that need to be up and running at the end of the exam.
5. Work on all storage-related tasks.
6. Create all required user and group accounts.
7. Set permissions.
8. Make sure SELinux is working and configured as required.
9. Work on everything else.

The third thing that you need to know about the exam is that you should reboot at least a couple of times. A successful reboot allows you to verify that everything is working up to the moment you have rebooted. Before rebooting, it is a good idea to remove the `rhgb` and `quiet` options from the GRUB boot loader. Removing them allows you to see what actually happens and makes troubleshooting a lot easier.

Do not reboot only at the end of the exam, because if at that moment you encounter an issue, you might not have enough time to fix it. You should at least make sure to reboot after working on all storage-related assignments.

The Nondisclosure Agreement

The RHCSA certification is the most sought-after certification that currently exists in IT. It represents a real value because it demonstrates that the person who has passed the RHCSA exam is able to work through a list of realistic assignments and complete the job successfully. It is in everybody's interest to help maintain this high value of the RHCSA certification. The nondisclosure agreement (NDA) is an important part of that.

The RHCSA exam requires demonstrating real skills because the content of the exam is not publicly available. Please help keep these exams valuable by not talking about questions that you have seen on the exam. Anyone who knows before the exam which questions will be asked will have an easier exam than you had, which means that the value of the certification will diminish and will make your effort less valuable. So, please help protect what you have worked so hard for and do not talk about exam content to anyone.

Also, you should realize that there is a penalty for disclosing exam questions after you have signed the NDA. You will lose your certification if you have passed the exam, or you will become barred from retaking it if you did not pass.

Theoretical Pre-Assessment Exam

This chapter provides an RHCSA theoretical pre-assessment exam to help you determine what you know and what you do not know. This theoretical exam is provided so that you can assess your skills and determine the best route forward for studying for the exam.

The RHCSA exam is a 100% practical exam. Therefore, you need to work on actual configuration tasks, and you must deliver a working configuration at the end of the exam. Therefore, passing this practical exam requires that you have a working knowledge of RHEL 8. This chapter helps you check whether you have the requisite knowledge.

In the following pre-exam theoretical exam, you are asked how you would approach some essential tasks. The purpose is to check for yourself whether you are on the right track. You do not have to provide a detailed step-by-step procedure. You just need to know what needs to be done. For instance, if the question asks how to set the appropriate SELinux context type on a nondefault web server document root, you know what you need to know if you say “check the `semanage-fcontext` man page.” If you do not have the answers to any of these questions, you know that you need to do additional studying on those topics.

In this theoretical pre-assessment exam, some key elements are covered. This test is *not* 100% comprehensive; it just focuses on some of the most essential skills.

1. You need to create a shared group environment where members of the group sales can easily share permissions with one another. Which approach would you suggest?
2. You need to change the hostname of the computer to something else and do it persistently. How would you do that?
3. On your disk, you have to create a logical volume with a size of 500 MB and the name `my_lv`. You do not have LVM volumes yet. List the steps to be taken to create the logical volume and mount it as an Ext4 file system on the `/data` directory. Also ensure that the extent size this logical volume uses is 8 MiB.

4. On the logical volume created in step 3, you need to set an ACL that gives members of the account group read and execute permissions. All other permission settings can be left as they are. How would you do this?
5. While booting, your server gives an error and shows “Enter root password for maintenance mode.” What is the most likely explanation for this problem?
6. You need to access a repository that is available on `ftp://server.example.com/pub/repofiles`. How would you do this?
7. You want to manually edit the network configuration by modifying the relevant configuration file for the `eth0` interface. What is the name of this file? Do you also need to do something else to make sure that the configuration is not changed back again automatically? Which service needs to be restarted to make the changes effective?
8. What configuration line would you add to which configuration file to schedule a cron job that executes the command **logger it is 2 AM** at 2 a.m. on every weekday? (You need to exclude Saturday and Sunday.)
9. You have configured your web server to listen at port 8082, and now it doesn’t start anymore. How do you troubleshoot?
10. You have access to the server console, but you do not have the root password to log in to that server. Describe step by step what you would do to get access to the server by changing the password of user root.
11. Describe exactly what you need to do to automatically mount home directories for LDAP users. The home directories are on `nfs://server.example.com/home/ldapusers`, and they should be automounted at `/home/ldapusers` on your local machine.
12. You need to install the RPM package that contains the file `sealert`, but you have no clue what the name of this package is. What is the best way to find the package name?
13. You have just downloaded a new kernel file from an FTP server; the update is not available in a repository. How do you use it to update your kernel in a way that allows you to install the new kernel but still keep the old kernel available for booting as a backup in case things go wrong?
14. You are trying to find relevant man pages that match the keyword `user`. You type `man -k user` but get the “nothing appropriate” message. How can you fix this?
15. How do you add a user to a new secondary group with the name `sales` without modifying the existing (secondary) group assignments?

16. How would you create a 5-GiB VDO volume with the name vdodata and mount it automatically on /vdodata?
17. How would you configure timesync, such that your server is synchronizing time with server10.example.com?
18. How would you check the recommended tuned profile and set it as the default profile?

This page intentionally left blank

RHCSA Practice Exam A

General Notes

Here are some tips to ensure your exam starts with a clean environment:

- You do not need external servers or resources.
 - Do *not* register or connect to external repositories.
 - Install a new VM according to the instructions in each practice exam.
 - No sample solutions are provided for these practice exams. On the real exam, you need to be able to verify the solutions for yourself as well.
 - You should be able to complete each exam within two hours.
1. Install a RHEL 8 or CentOS 8 virtual machine that meets the following requirements:
 - 2 GB of RAM
 - 20 GB of disk space using default partitioning
 - One additional 20-GB disk that does not have any partitions installed
 - Server with GUI installation pattern
 2. Create user **student** with password **password**, and user **root** with password **password**.
 3. Configure your system to automatically loop-mount the ISO of the installation disk on the directory **/repo**. Configure your system to remove this loop-mounted ISO as the only repository that is used for installation. Do *not* register your system with **subscription-manager**, and remove all reference to external repositories that may already exist.
 4. Reboot your server. Assume that you don't know the root password, and use the appropriate mode to enter a root shell that doesn't require a password. Set the root password to **mypassword**.
 5. Set default values for new users. Set the default password validity to 90 days, and set the first UID that is used for new users to 2000.

6. Create users **edwin** and **santos** and make them members of the group **livingopensource** as a secondary group membership. Also, create users **serene** and **alex** and make them members of the group **operations** as a secondary group.
7. Create shared group directories **/groups/livingopensource** and **/groups/operations**, and make sure the groups meet the following requirements:
 - Members of the group **livingopensource** have full access to their directory.
 - Members of the group **operations** have full access to their directory.
 - New files that are created in the group directory are group owned by the group owner of the parent directory.
 - Others have no access to the group directories.
 - Members of the group **operations** have read access to the directory **/groups/livingopensource**.
8. Create a 2-GiB volume group, using 8-MiB physical extents. In this volume group, create a 500-MiB logical volume with the name **mydata**, and mount it persistently on the directory **/mydata**.
9. Find all files that are owned by user **edwin** and copy them to the directory **/rootedwinfiles**.
10. Schedule a task that runs the command **touch /etc/motd** every day from Monday through Friday at 2 a.m.
11. Add a new 10 GiB virtual disk to your virtual machine. On this disk, add a VDO volume with a size of 20 GiB and mount it persistently.
12. Create user **bob** and set this user's shell so that only this user can change the password.
13. Install the **vsftpd** service and ensure that it is started automatically at reboot.

RHCSA Practice Exam B

General Notes

Here are some tips to ensure your exam starts with a clean environment:

- You do not need external servers or resources.
 - Do not register or connect to external repositories.
 - Install a new VM according to the instructions in each practice exam.
 - No sample solutions are provided for these practice exams. On the real exam, you need to be able to verify the solutions for yourself as well.
 - You should be able to complete each exam within two hours.
1. Install a RHEL 8 or CentOS 8 virtual machine that meets the following requirements:
 - 2 GB of RAM
 - 20 GB of disk space using default partitioning
 - One additional 20-GB disk that does not have partitions installed
 - Server with GUI installation pattern
 2. Create user **student** with password **password**, and user **root** with password **password**.
 3. Configure your system to automatically loop-mount the ISO of the installation disk on the directory **/repo**. Configure your system to remove this loop-mounted ISO as the only repository that is used for installation. Do *not* register your system with **subscription-manager**, and remove all reference to external repositories that may already exist.
 4. Create a 1-GB XFS partition on **/dev/sdb**. Mount it persistently on the directory **/mydata**, using the label **mylabel**.
 5. Set default values for new users. Ensure that an empty file with the name **NEWFILE** is copied to the home directory of each new user that is created.

6. Create users **laura** and **linda** and make them members of the group **livingopensource** as a secondary group membership. Also, create users **lisa** and **lori** and make them members of the group **operations** as a secondary group.
7. Create shared group directories **/groups/livingopensource** and **/groups/operations** and make sure these groups meet the following requirements:
 - Members of the group **livingopensource** have full access to their directory.
 - Members of the group **operations** have full access to their directory.
 - Users should be allowed to delete only their own files.
 - Others should have no access to any of the directories.
8. Create a 2-GiB swap partition and mount it persistently.
9. Resize the LVM logical volume that contains the root file system and add 1 GiB.
10. Set your server to use the recommended tuned profile.
11. Create user **vicky** with the custom UID **2008**.
12. Configure your server to synchronize time with **myserver.example.com**. (Note that this server does not have to exist.)
13. Install a web server and ensure that it is started automatically.

Answers to the “Do I Know This Already?” Quizzes and Review Questions

Answers to the “Do I Know This Already?” Quizzes

Chapter 1

1. A. Fedora is an experimental/enthusiast version containing many components that may or may not make it into the RHEL distribution tree and onto the RHCSA exam.
2. D. All RHEL software updates are made available in CentOS as well.
3. A. In particular, when working with virtual machines, you’ll be happy to have a GUI at your disposal.
4. C. XFS is used as the default file system. When Red Hat decided which file system to use as the default file system, Btrfs was not stable enough yet.
5. A. The size of an XFS file system cannot be reduced.
6. C. The Fedora project tries to make a stable distribution as well. There are many Fedora users around the globe who use it as a production distribution.
7. D. The Troubleshoot an Existing Installation option is available when booting from disk, not on the Installation Summary screen.
8. D. You are allowed to use an unsecure password; you just have to confirm it twice.
9. D. Language settings can be changed after installation. This is done easily through the Settings option in the graphical interface.
10. B. Even if it makes sense having /var on a dedicated partition, this is not part of a default installation.

Chapter 2

1. B. You first must redirect the standard output to a file, and then **2>&1** is used to specify that errors are to be treated the same way.
2. A and C. `/etc/profile` is the file that is processed for all users who are starting a login shell. A user-specific version exists as well, with the name `~/.bash_profile`.
3. C. On Linux, the current directory is not set in the `PATH` variable.
4. D. A pipe is used to process the output of the first command and use it as input of the second command.
5. A. The command **history -c** removes the in-memory state from the history file of current history. Remove `~/.bash_history` also to make sure that all history is removed. As an alternative to removing the `.bash_history` file, the user can type **history -c**, followed by **history -w**.
6. D. Ctrl-X is not a valid history command.
7. D. Bash completion works for commands, files, variables and other names if configuration for that has been added (like hostnames for the SSH command).
8. A. You need the command `:%s/old/new/g` to replace all instances of *old* with *new*. `%` means that it must be applied on the entire file. `s` stands for substitute. The `g` option is used to apply the command to not only the first occurrence in a line (which is the default behavior) but all occurrences in the line.
9. B. The `/etc/motd` file contains messages that are displayed after user login on a terminal session. If you want to show a message before users log in, edit the `/etc/issue` file.
10. C. The **man -k** command uses a database to find the keywords you are looking for. On RHEL 8, this database is updated with the **mandb** command. On older versions of RHEL (prior to RHEL 7), the **makewhatis** command was used instead.

Chapter 3

1. D. Program files that are not required to boot a system are typically stored in a subdirectory below the `/usr` directory. In old versions of RHEL, essential binaries and system binaries were stored in `/bin` and `/sbin`, but on modern versions of RHEL, these directories are a symbolic link to `/usr/bin` and `/usr/sbin`.
2. C. The `/var` directory is used on Linux to store files that may grow unexpectedly.
3. A. The `/etc/` directory contains configuration files that are needed while your server boots. Putting `/etc` on a dedicated device would make your server unbootable.

4. C. The **df -h** command shows mounted devices and the amount of disk space currently in use on these devices. The **-T** option helps in recognizing real file systems (as opposed to kernel interfaces) because it shows the file system type as well.
5. C. The option **-a** shows hidden files, **-l** gives a long listing, **-t** sorts on modification time, which by default shows newest files first, and **-r** reverts the sorting so that newest files are shown last.
6. C. To copy hidden files as well as regular files, you need to put a **.** after the name of the directory the files are in. Answer A copies hidden files as well, but it creates a subdirectory **\$USER** in the current directory.
7. A. The **mv** command enables you to move files and rename files.
8. D. In hard links, no difference exists between the first hard link and subsequent hard links.
9. C. The option **-s** is used to create a symbolic link. While creating a link, you first have to specify the source, and next you specify the destination.
10. C. Use the option **-r** to add one single file to an archive you have created with **tar**.

Chapter 4

1. A. The **head** command by default shows the first ten lines in a text file.
2. D. The **wc** command shows the number of lines, words, and characters in a file.
3. D. When using **less**, the G key brings you to the end of the current file.
4. A. The **-d** option is used to specify the field delimiter that needs to be used to distinguish different fields in files while using **cut**.
5. A. The **sort** command can sort files or command output based on specific keys. If no specific key is mentioned, sorting happens based on fields. The option **-k3** will therefore sort the third field in the output of the **ps aux** command.
6. D. When used in a regular expression, the **^** sign in front of the text you are looking for indicates that the text has to be at the beginning of the line.
7. A. The **?** regular expression is used to refer to zero or one of the previous characters. This makes the previous character optional, which can be useful. If the regular expression is **colou?r**, for example, you would get a match on *color* as well as *colour*.
8. B. The **.** is used as a regular expression to refer to any single character.

9. D. The **awk** command first needs to know which field separator should be used. This is specified with the **-F :** option. Then, it needs to specify a string that it should look for, which is **/user/**. To indicate that the fourth field of a matching file should be printed, you need to include the **{ print \$4 }** command.
10. B. Use **grep -v** to exclude from the results lines containing the regular expression.

Chapter 5

1. B. The console is the screen you are working from. On the console, a terminal is started as the working environment. In the terminal, a shell is operational to interpret the commands you are typing.
2. A. The console is the screen you are working from. On the console, a terminal is started as the working environment. In the terminal, a shell is operational to interpret the commands you are typing.
3. C. The console is the screen you are working from. On the console, a terminal is started as the working environment. In the terminal, a shell is operational to interpret the commands you are typing.
4. B. The six virtual consoles that are available on Linux by default are numbered **/dev/tty1** through **/dev/tty6**. The device **/dev/pts/6** is used to refer to the sixth pseudo terminal, which is created by opening six terminal windows in a graphical environment.
5. A and C. A pseudo terminal device is created when opening new terminals using SSH or from the graphical interface.
6. D. Typically, a server reboot is only necessary after making changes to the kernel and kernel modules that are in use. Changing the network configuration does not normally require a reboot, because it is possible to just restart the network service.
7. C. Windows has no native support for SSH. You need to install PuTTY or similar software to remotely connect to Linux using SSH.
8. D. Key fingerprints of hosts that you have previously connected to are stored in your home directory, in the subdirectory **.ssh** in a file with the name **known_hosts**.
9. C. The **ForwardX11** option in the **/etc/ssh/ssh_config** file enables support for graphical applications through SSH.
10. C. To initiate key-based remote authentication, you should copy the public key to the remote server. The most convenient way to do so is using the **ssh-copy-id** command.

Chapter 6

1. C. Privileged users are the opposite of unprivileged users. A privileged user can execute tasks in kernel space, without further restriction. By default, only the user root exists as a privileged user.
2. D. In the sudo configuration file, all members of the group wheel by default get access to all administrator tasks.
3. B. The **runas** command does not exist on Linux.
4. B. The hashed user passwords are stored in /etc/shadow.
5. C. The file /etc/default/useradd is read for default settings when new user accounts are created.
6. A. The **chage -l** command enables you to manage password properties.
7. B. There is no file /etc/.profile.
8. A. The **vigr** command creates a copy of the /etc/group file so that changes can be applied safely.
9. C and D. The **id** and **groups** commands show a list of all groups a user is a member of.
10. D. When a file with the name /etc/nologin.txt is created, only the root user is allowed to log in, and nobody else. In this file you can specify a message that will be shown to users who are trying to log in.

Chapter 7

1. C. The **newgrp** command is used to set the effective primary group, which will affect default group ownership on new files until the current shell session is ended. The **chgrp** command is used to set the group owner of an existing file; **chgrp** is not related to any user account, and it affects newly created files only.
2. A. The **find / -user linda** command searches all files that are owned by user linda. Notice that **find** also has a **-uid** option that allows you to locate files based on a specific UID setting. This does not allow you to search files based on a username, but it will let you find files based on the UID of a specific user.
3. C. **chgrp myfile sales** does not set group ownership for the file myfile. The order in this command is wrong; **chgrp** first needs the name of the group, followed by the name of the owner that needs to be set.
4. C. When used in relative mode, the three digits are used to specify user, group, and others permissions. The value 6 is used to apply read and write.

5. C. The **chmod g+s /dir** command adds (+) the SGID permission to /dir. **chmod u+s /dir** adds SUID to the directory, **chmod g-s /dir** removes the SGID permission, and the 1 in **chmod 1770 /dir** would set sticky bit and not SGID.
6. D. ACL support is not offered by default on all file systems. If you get an “operation not supported” error message, make sure to add the **acl mount** option and remount the file system. Note that this is not required on RHEL 8 XFS, where ACL support is set by default.
7. B. Although answers A and B will both set default ACLs, answer B is better because it adds x to the permissions. Without x, members of the group sales will have no way to enter the directory using the **cd** command.
8. A. The umask is a system-wide setting and cannot be used to apply to specific directories only. Use a default ACL as shown in answer A to perform this task.
9. C. In a umask, 0 in the first position gives all permissions to the file owner, 2 in the second position ensures that members of the group owner can read files, and 7 in the third position takes away all permissions for others.
10. C. The **lsattr** command shows current attribute settings to files. The **ls** command is not capable of showing file attributes, and the other commands that are listed do not exist.

Chapter 8

1. D. Based on the /26 subnet mask, the networks are 192.168.4.0, 192.168.4.64, 192.168.4.128, and 192.168.4.192. That means that IP addresses II, III, and IV belong to the same network.
2. B. The 169.254.0.0 network address does not belong to the private address ranges, which are 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16. The address 169.254.0.0 is from the APIPA range. This is a range of IP addresses that can be automatically self-assigned by a client that is trying to reach an unreachable DHCP server.
3. C. By default, the network device names are based on the device type, location, and identification. You’ll find the legacy name eth0 only if the driver of the network card is not capable of revealing information about its physical location.
4. D. Use of the **ifconfig** command is deprecated; use the **ip** command instead. The **ip addr show** command shows information about the state of the interface as well as the current IP address assignment.

5. A. The network service no longer exists in RHEL 8.
6. B. The `nmcli-examples` man page was created to make working with the long commands in `nmcli` a bit easier. Note that **nmcli** also has excellent command-line completion features.
7. C. On RHEL 8, `nmtui` is the default utility to set and modify the network interface. Note that RHEL 8 no longer has the system-config utilities for configuring networking that were available in older versions of RHEL.
8. D. When the connection is added, you use **ip4** and **gw4** (without a *v*).
9. A. You should not set the DNS servers directly in `/etc/resolv.conf`, because that file is automatically written by the NetworkManager service.
10. C. The name of the configuration file that contains the hostname is `/etc/hostname`. You should use **hostnamectl** to change its contents.

Chapter 9

1. D. The `gpgcheck=` line indicates whether to check the integrity of packages in the repository using a GPG key. Although useful, this is not mandatory.
2. B. If a RHEL system is not registered with RHEL, no repositories are available. This is important to realize for the RHCSA exam, because it means that you need to connect to a repository before you can install anything.
3. C. Use `baseurl` to specify which URL to use. If the URL is based on the local file system, it uses the URI `file://` followed by the path on the local file system, which in this case is `/repo`. This explains why there are three slashes in the `baseurl`.
4. D. GPG package signing is used to set a checksum on packages, so that altered packages can easily be recognized. The main purpose of signing packages is to make it easy to protect packages on Internet repositories. For internal repositories that cannot be accessed by Internet users, the need to add GPG package signatures is less urgent.
5. D. Both the commands **yum provides** and **yum whatprovides** can be used to search for files within a specific package. The file pattern must be specified as `*/filename` or as a full path, which is why answer D is the only correct answer. Without the `*/` in front of the filename, you may get a match that is based on the package description, not on the filename.
6. B. The `yum` module application stream allows for working with different versions of user space software side by side.

7. C. To install a specific profile from a yum module application stream, add the profile name to the application stream version using a `/`.
8. A. The **yum install** command installs individually downloaded RPM files while looking for package dependencies in the current repositories. This is better than using **rpm -ivh**, which does not consider the yum repositories. In earlier versions of RHEL, the **yum localinstall** command was used to install packages that were downloaded to the local file system, but this command is now deprecated.
9. C. Use the **rpm -qf** command to find which RPM package a specific file comes from.
10. C. The **--scripts** option checks whether there are scripts in an RPM package. If you want to query the package file and not the database of installed RPMs, you need to add the **-p** option to the **-q** option, which is used to perform RPM queries.

Chapter 10

1. B and D. There are two different types of processes that each request a different management approach. These are shell jobs and daemons. A cron job and a thread are subdivisions of these generic categories.
2. B. The Ctrl-Z command temporarily freezes a current job, after which the **bg** command can be used to resume that job in the background.
3. A. The Ctrl-C command cancels the current job. Ctrl-D sends the EOF character to the current job, which can result in a stop if this allows the job to complete properly. The difference with Ctrl-C is that the job is canceled with no regard to what it was doing. The Ctrl-Z keystroke freezes the job.
4. A and B. Individual threads cannot be managed by an administrator. Using threads makes working in a multi-CPU environment more efficient because one process cannot be running on multiple CPUs simultaneously, unless the process is using threads.
5. A. The **ps ef** command shows all processes, including the exact command that was used to start them.
6. C. To increase process priority, you need a negative nice value. -20 is the lowest value that can be used.
7. C. Use the **renice** command to change priority for currently running processes. To refer to the process you want to renice, use the **-p** option.
8. B. **kill** is not a current command to send signals to processes.

9. A. To change the process priority from **top**, use **r** for renice.
10. B. To set the tuned performance profile, use **tuned-adm profile**, followed by the name of the profile you want to set.

Chapter 11

1. A. The **--type=service** argument shows all currently loaded services only.
2. C. Wants are specific to a particular system and for that reason are managed through `/etc/systemd/system` (not `/usr/lib/systemd/system`).
3. D. Masking a service makes it impossible to enable it.
4. D. Running(dead) is not a valid status for systemd services. You see (dead) for units that currently are not active.
5. D. Socket units monitor socket activity, which may consist of a file being accessed or a network port being accessed. They do not monitor PATH activity. This is done by the path unit type.
6. B. udev is not a valid systemd unit type. All others are.
7. B. Answers A and B are very similar, but answer A uses the wrong command. You have to use the **systemctl** command, not the **systemd** command.
8. D. The `SYSTEMD_EDITOR` variable defines which editor to use. You need to set a full path name for this editor.
9. D. The Wants dependency type defines that a specific unit is wanted, without setting a hard requirement.
10. A. The word order is wrong. It should be **systemctl start unit**, not **systemctl unit start**.

Chapter 12

1. B. The **systemctl** command on RHEL 8 is used to manage services. If used with the **status** option, the current state of the service is checked, and recent log messages about the service are displayed as well. The name of the service to check in this case is `crond`.
2. C. The fields in cron timing are minute, hour, day of month, month, and day of week. Answer C matches this pattern to run the task on the seventh day of the week at 11 a.m.

3. A. To launch a job from Monday through Friday, you should use **1-5** in the last part of the time indicator. The minute indicator ***/5** will launch the job every five minutes.
4. A and D. You cannot modify user cron files directly, but have to go through the crontab editor. This editor is started with the **crontab -e** command.
5. B. The **/etc/cron.d** directory is used to store cron files for individual services that need tasks to be executed through cron. This directory is mostly filled by installing RPM files that contain cron jobs.
6. A. Although cron jobs that are added to **/etc/crontab** will be executed, the **/etc/crontab** file is considered a system file that should not be modified directly.
7. D. anacron is a service that takes care of executing jobs on a regular basis where it is not necessary to specify a specific time.
8. D. By default, the **cron.allow** file does not exist. If it exists, a user must be listed in it in order to program cron jobs.
9. B. The Ctrl-D key sequence sends the end-of-file (EOF) character to the at shell and closes it.
10. C. The **atq** command queries the at service and provides an overview of jobs currently scheduled for execution.

Chapter 13

1. C. journald is not a replacement of rsyslogd. It is an additional service that logs information to the journal. In RHEL 8, they are integrated to work together to provide you with the logging information you need.
2. D. Most messages are written to the **/var/log/messages** file, but authentication-related messages are written to **/var/log/secure**. Check the contents of **/etc/rsyslog.conf** and look for **authpriv** to find out what exactly is happening for this facility.
3. C. SELinux events are logged through the audit service, which maintains its log in **/var/log/audit/audit.log**.
4. A. The rsyslogd configuration file is **/etc/rsyslog.conf**.
5. D. The **/etc/sysconfig/rsyslog** file is the default location to change rsyslogd startup parameters by editing the **SYSLOGD_OPTIONS** parameter.
6. C. rsyslogd destinations often are files. For further processing, however, log information can be sent to an rsyslogd module. If this is the case, the name of the module is referred to as **:modulename:**.

7. D. The local facilities local0 through local7 can be used to configure services that do not use rsyslog by default to send messages to a specific rsyslog destination, which needs to be further configured in the rsyslog.conf file.
8. A. The logrotate service can rotate files based on the maximal file size. To configure this, the recommended way is to drop a file in /etc/logrotate.d containing parameters for this specific file.
9. D. If systemd-journald has been configured for persistent storage of the journal, the journal is stored in /run/log/journal. Note that by default the systemd journal is not stored persistently.
10. A. To make the journald journal persistent, you have to create a directory /var/log/journal and set the appropriate permissions to that directory.

Chapter 14

1. A. In GPT, there is no longer a need to differentiate between primary, extended, and logical partitions; in fact, it is not even possible. Using logical partitions is not an advantage due to the limited number of primary partitions available on MBR disks.
2. B. 1 pebibyte (PiB) is $1024 \times 1024 \times 1024 \times 1024 \times 1024$ bytes.
3. C. Partition type 83 is normally used to create Linux partitions.
4. C. KVM virtual machines use the virtio driver to address hard disks. This driver generates the device /dev/vda as the first disk device.
5. C. A disk can have one partition table only. For that reason, it is not possible to have MBR and GPT partitions on the same disk.
6. B. XFS is used as the default file system; partitions can still be formatted with other file systems, like Ext4.
7. D. The **blkid** command shows all file systems, their UUID, and if applicable, their label.
8. D. To mount a file system based on its UUID, use UUID=nnnn in the /etc/fstab device column.
9. B and D. To check a file system upon boot, but only after the root file system has been checked successfully, put a 2 in the sixth column in /etc/fstab.
10. B. The **_netdev** mount option is used to specify that the file system depends on the network to be present before it can be mounted.

Chapter 15

1. B. It is common to create a file system on top of a logical volume, but this is not a requirement. For instance, a logical volume can be used as a device that is presented as a disk device for a virtual machine.
2. C. Copy on write is a feature that is offered by modern file systems, such as Btrfs. It copies the original blocks a file was using before creating a new file, which allows users to easily revert to a previous state of the file. Copy on write is not an LVM feature.
3. D. On a GPT disk, LVM partitions must be flagged with the partition type 8e00.
4. C. The **lvcreate** command is used to create logical volumes. Use **-n name** to specify the name. The option **-l 50%FREE** will assign 50% of available disk space, and **vgdata** is the volume group it will be assigned from.
5. B and C. The **pvdiskdisplay** command is used to show extensive information about physical volumes. The **pvs** command shows a summary of essential physical volume properties only.
6. C. You can increase the size of an XFS file system, but it cannot be decreased.
7. A. To write metadata, each Stratis volume requires 527 MiB of storage.
8. C. You need to mount Stratis volumes using the UUID and not the device name. Also, the option **x-systemd.requires=stratisd.service** needs to be included to ensure that the stratisd.service is loaded before systemd tries to mount the Stratis volume.
9. A, C, D. Because of its data deduplication features, VDO is doing very well in environments where redundant data is written. This particularly applies to cloud storage, containers, and virtualization.
10. B. The **-K** option is used with the **mkfs.xfs** command to ensure that empty data blocks are not discarded immediately.

Chapter 16

1. A. A tainted kernel is caused by drivers that are not available as open source drivers. Using these may have an impact on the stability of the Linux operating system, which is why it is good to have an option to recognize them easily.
2. B. The **dmesg** utility shows the contents of the kernel ring buffer. This is the area of memory where the Linux kernel logs information to, so it gives a clear overview of recent kernel events.

3. A. The **uname -r** command shows the current kernel version. The **uname -v** command gives information about the hardware in your computer, and the **procinfo** command does not exist.
4. C. The `/etc/redhat-release` version contains information about the current version of RHEL you are using, including the update level.
5. A. On a systemd-based operating system such as RHEL 8, the `systemd-udev` process takes care of initializing new hardware devices.
6. B. Default rules for hardware initialization are in the directory `/usr/lib/udev/rules.d`; custom rules should be stored in `/etc/udev/rules.d`.
7. C. The **modprobe** command is the only command that should be used for managing kernel modules, as it considers kernel module dependencies as well. Use **modprobe** to load a kernel module and **modprobe -r** to unload it from memory.
8. C. The **lspci -k** command lists devices that are detected on the PCI bus and supporting kernel modules that have been loaded for those devices. Alternatively, **lspci -v** shows more verbose information about modules that are currently loaded.
9. C. The `/etc/modprobe.d` directory is used for files that create custom configurations. The files `/etc/modules.conf` and `modprobe.conf` were used for this purpose in the past. On RHEL 8, kernel module parameters are passed through `/usr/lib/modprobe.d` if they are used for operating system–managed permanent parameters.
10. C and D. Kernels are not updated; they are installed. You can use either **yum update kernel** or **yum install kernel** to do so. There are no additional requirements, which makes answers C and D false.

Chapter 17

1. A. The **systemctl enable** command creates a want for the current unit in the target that is listed in the `[Install]` section in the service unit file.
2. C. Servers typically don’t run a graphical interface and will start the `multi-user.target`.
3. B. There is no `restart.target`.
4. D. Unit files contain an `[Install]` section that is used to specify in which target the unit should be started.

5. A. The required statement is **AllowIsolate**. All other statements mentioned here are invalid.
6. A. To switch from a target with more unit files to a target with fewer unit files, use **systemctl isolate**.
7. B. The multi-user.target corresponds roughly to runlevel 3 as used in a System V environment.
8. D. Changes to GRUB 2 need to be applied to /etc/default/grub, not to /boot/grub2/grub.cfg. The /boot/grub2/grub.cfg file cannot be edited directly; you have to apply changes to /etc/default/grub and run the **grub2-mkconfig** command to write them to the appropriate configuration file.
9. B. The **grub2-mkconfig** command enables you to regenerate the GRUB 2 configuration. The result, by default, is echoed to the screen. Use redirection to write it to a file.
10. B. The /boot/efi/EFI/redhat/grub.cfg file is used to store the GRUB 2 bootloader on a UEFI system.

Chapter 18

1. C. During the boot procedure, the GRUB 2 boot loader gets loaded first. From here, the kernel with the associated initramfs is loaded. Once that has completed, systemd can be loaded.
2. B. The Ctrl-X key sequence leaves the GRUB 2 shell and continues booting.
3. B. The /etc/dracut.conf file is used for managing the initramfs file system.
4. D. The **rd.break** boot option enters at the end of the initrd phase. The root file system has not been mounted on / yet, which allows for easy troubleshooting.
5. A and C. The **rhgb** and **quiet** boot options make it impossible to see what is happening while booting.
6. B. The emergency.target systemd target gives just a root shell and not much more than that. All other options that are mentioned also include the loading of several systemd unit files.
7. C. If you do not get to a GRUB 2 boot prompt, you cannot select an alternate startup mechanism. This situation requires you to use a rescue disk so that GRUB can be reinstalled. If the kernel or initramfs cannot load successfully, you might need to use a rescue disk also, but in many cases an alternate kernel is provided by default.

8. C. The **mount -o remount,rw /** option remounts the / file system in read/write mode.
9. A. Because the error occurs before the GRUB 2 menu is loaded, the only option to fix this is to use a rescue disk.
10. C. The **rd.break** method is applied before SELinux is active, which removes the SELinux context label. You’ll have to set this label again before rebooting.

Chapter 19

1. B. The first line of a Bash shell script contains the *shebang*, which defines the subshell that should be used for executing the script code.
2. A. The **exit 0** statement at the end of a script is an optional one to inform the parent shell that the script code was executed successfully.
3. C. The **read** statement stops a script, which allows a user to provide input. If **read** is used with a variable name as its argument, the user input is stored in this variable.
4. D. The first argument is referred to as \$1. To store \$1 in a variable with the name NAME, use the command **NAME=\$1**. Make sure that no spaces are included, which is why answer A is incorrect.
5. D. Both \$@ and \$* can be used to refer to all arguments that were provided when starting a script, but \$@ is the preferred method because it enables the script to distinguish between the different individual arguments, whereas \$* refers to all the provided arguments as one entity.
6. D. A conditional loop that is started with **if** is closed with **fi**.
7. C. If within an **if** loop a new conditional check is opened, this conditional check is started with **elif**.
8. B. After stating the condition in a **for** loop, **do** is used to start the commands that need to be started when the condition is true.
9. D. The **mail** command needs its subject specified with the **-s** option. The **mail** command normally waits until a dot is entered on an empty line to start sending the message. This dot can be fed to the **mail** command using STDIN redirection, using **<**.
10. A. In a **case** statement, the different options are proposed with a) behind them. *) refers to all other options (not specifically specified in the script).

Chapter 20

1. A. X11 forwarding applies to sessions that have already been authorized. Disabling it does not protect against brute-force attacks.
2. C. The **AllowUsers** parameter can be used to restrict SSH access to specific users only.
3. D. To change the port on which SSH is listening in SELinux, the port must be allowed as well. To do so, use **semanage port**. The context type needs to be set to **ssh_port_t**.
4. B. The **MaxAuthTries** setting starts logging failed login attempts after half the number of attempts specified here.
5. D. Login-related settings are logged to **/var/log/secure**.
6. C. SSH is trying to do a reversed lookup of the DNS name belonging to a target IP address. If faulty DNS configuration is used, this will take a long time.
7. D. The **UseDNS** option has nothing to do with SSH session keepalive.
8. A. SSH client settings that apply to a specific user only can be stored in **~/.ssh/ssh_config**.
9. C. The **ssh-agent** command adds an SSH credentials cache to a shell. Next, you need to run **ssh-add** to add a specific key to the cache.
10. A. By default, an SSH server can support ten sessions only.

Chapter 21

1. A. The **httpd** package contains the core components of the Apache web server. It can be installed using **yum install httpd**.
2. A. The default Apache configuration file is **/etc/httpd/conf/httpd.conf**.
3. C. The **DocumentRoot** parameter specifies where the Apache web server will look for its contents.
4. A. The **ServerRoot** parameter defines where Apache will look for its configuration files. All file references in the **httpd.conf** configuration file are relative to this directory.
5. B. The **/etc/httpd/conf** directory contains the main Apache configuration file **httpd.conf**.

6. D. The `/etc/httpd/modules.d` directory contains configuration files that are used by specific Apache modules.
7. C. The `/etc/httpd/conf.d` directory is used by RPMs that can drop files in that directory without changing the contents of the main Apache configuration file.
8. A. The name-based virtual host is used as the default virtual host type. It allows multiple virtual hosts to be hosted on the same IP address.
9. A. The **VirtualHost** parameter is used to open a virtual host definition. `*` refers to all IP addresses, and `:80` defines the port it should listen on.
10. C. No additional packages need to be installed to enable virtual hosts. Virtual hosts are supported through the default `httpd` RPM package.

Chapter 22

1. D. Enabled is not a valid mode that can be set using **setenforce** or the `/etc/sysconfig/selinux` configuration file.
2. A and D. The **getenforce** command is used to request the current SELinux mode. The **sestatus** command can be used also. It shows the current mode, and some additional security-related information as well.
3. A. For basic SELinux configuration, you need to make sure that the appropriate context type is set. User and role are for advanced use only.
4. D. SELinux security can be applied to users, files, and ports.
5. C. The **-Z** option displays SELinux-related information and can be used with many commands.
6. D. This is the only command that provides correct usage information about **semanage**. Remember that **chcon** should be avoided at all times.
7. B. When moving a file, the original file context is moved with the file. To ensure that the file has the context that is appropriate for the new file location, you should use **restorecon** on it.
8. B. To change Booleans, use **setsebool**; to make the change persistent, use **-P**.
9. A. SELinux messages are logged by `auditd`, which writes the log messages to `/var/log/audit/audit.log`. Only if **sealert** is installed are messages written to `/var/log/messages` as well, but that does not happen by default.
10. D. SELinux log messages always contain the text *avc*, which stands for access vector cache.

Chapter 23

1. A. On a default configuration, there is no untrusted zone in firewalld.
2. C. Netfilter is the name of the firewall implementation in the Linux kernel. Different toolsets exist to manage netfilter firewalls. Iptables has been the default management interface for a long time, and in Red Hat Enterprise Linux 7, firewalld was added as an alternative solution to manage firewalls.
3. D. Firewalld and iptables are mutually exclusive.
4. C. The **firewall-cmd --get-services** command shows all services that are available in firewalld.
5. C. The name of the GUI tool that can be used to manage firewall configurations is **firewall-config**.
6. B. Answer B shows the correct syntax.
7. A. The trusted zone is provided for interfaces that need minimal protection.
8. D. Configuration that is added with the **--permanent** option is not activated immediately and needs either a restart of the firewalld service or the command **firewall-cmd --reload**.
9. B. The **--list-all** command without further options shows all configurations for all zones.
10. A. When working with **firewall-config**, you need to choose between the runtime configuration and the persistent configuration.

Chapter 24

1. B. In the default configuration, NFS share access is based on UID matching between the client and server. To enable anonymous user access, you need to specify the **sec=none** mount option.
2. A. NFSv4 does not offer straight integration with Active Directory. Similar functionality is provided by the option to use Kerberized security.
3. D. The nfs-utils package contains all that is needed to mount NFS shares.
4. D. **showmount** is using the NFS portmapper, which is using random UDP ports to make the connection. Portmapper traffic is not automatically allowed when the nfs service is added to the firewall because RPC ports that are needed by **showmount** are blocked by the firewall. Ensure that the rpc-bind and mountd services are added to the firewall to enable the **showmount** command to work as it should.

5. A. To authenticate to a Samba share, you need to use the **-o username=sambauser** option to specify the username.
6. D. To avoid having to put a username and password in clear text in the `/etc/fstab` file, you can use a credentials file while mounting the share.
7. D. You do not have to set permissions on the local file system for automount to be effective.
8. B. Each automounted directory should have a configuration file that has a name that matches the name of the automounted directory. For `/myfiles`, that would be `/etc/auto.myfiles`.
9. C. The first element is `*`, which refers to all directories that may be accessed in the local directory. The `&` matches this directory on the NFS share. The **-rw** option is used to specify NFS mount options.
10. A. Automount uses the `autofs` service.

Chapter 25

1. C. When booting, a server reads the hardware clock and sets the local time according to hardware time.
2. D. Hardware time on Linux servers typically is set to UTC, but local administrators may choose to make an exception to that general habit.
3. D. The **timedatectl** command, introduced as a new solution in RHEL 7, allows you to manage many aspects of time.
4. C. Atomic clocks can be used as a very accurate alternative to the normal hardware clock.
5. D. The `/etc/chrony.conf` file contains the default list of NTP servers that should be contacted on RHEL 8.
6. C. The **-s** option is used to set the current time, and to do so, military time format is the default.
7. A. To translate epoch time into human time, you need to put `@` in front of the epoch time string.
8. C. The **hwclock -c** command opens an interface that is refreshed every 10 seconds and shows the current hardware time, the system time, and the difference between the two of them.
9. D. When used without arguments, **timedatectl** gives a complete overview of current time settings on your server.

10. B. The **chronyc sources** command will show all current synchronization sources, as well as the stratum that was obtained from these sources.

Answers to the Review Questions

Chapter 1

1. Use CentOS to create an environment to practice for the exam without purchasing a RHEL license.
2. 32-bit RHEL does not support virtualization.
3. You need 512 MB of RAM to install a minimal system.
4. By default, updates and installation of additional software packages require Internet connectivity.
5. Use an ISO image to install a virtual machine on the computer.
6. To manage virtualization in an easy way, you need virt-manager, which is a GUI utility.
7. XFS is the default file system on RHEL 8.
8. You can install RHEL if you do not have Internet access. But you cannot register with RHN, so you will not have access to repositories after the installation has finished.
9. Repository access is the most important feature offered through RHN.
10. Use the Minimal Install pattern if you have a very limited amount of disk space available.

Chapter 2

1. A variable is a placeholder that contains a specific value and that can be used in scripts to work with dynamic contents.
2. **man -k** enables you to find the correct man page based on keyword usage.
3. Change /etc/bashrc to ensure a variable is set for every shell that is started.
4. Use **pinfo** to read the information.
5. Bash stores its history in ~/.bash_history.
6. **mandb** updates the database that contains man keywords.
7. Use **u** to undo the last modification you have applied in **vim**.

8. Add **2> /dev/null** to a command to ensure that it doesn’t show an error message.
9. Use **echo \$PATH** to read the current contents of the **\$PATH** variable.
10. Press Ctrl-r and type **dog**.

Chapter 3

1. **/etc** contains configuration files.
2. **ls -alt** displays a list of current directory contents, with the newest files listed first. (**-a** also shows files that have a name that starts with a dot.)
3. **mv myfile yourfile** renames **myfile** to **yourfile**.
4. **rm -rf /directory** wipes an entire directory structure, including all of its contents.
5. **ln -s /tmp ~** creates in your home directory a link to **/tmp**.
6. **cp /etc/[abc]* .** copies all files that have a name that starts with **a**, **b**, or **c** from the directory **/etc** to your current directory.
7. **ln -s /etc ~** creates in your home directory a link to **/etc**.
8. Use **rm symlink** to safely remove a symbolic link to a directory. If **rm** is aliased to **rm -i** and you do not want to answer yes for every individual file, use **\rm** instead.
9. **tar zcvf /tmp/etchome.tgz /etc /home** creates a compressed archive of **/etc** and **/home** and writes it to **/tmp/etchome.tgz**.
10. **tar xvf /tmp/etchome.tgz /etc/passwd** extracts **/etc/passwd** from **/tmp/etchome.tgz**.

Chapter 4

1. **ps aux | less** shows the results of **ps aux** in a way that is easily browsable.
2. **tail -n 5 ~/samplefile** shows the last five lines from **~/samplefile**.
3. **wc ~/samplefile**. You might use **-w** to show only the number of words.
4. Press Ctrl-C to stop showing output.
5. **grep -v -e '^#' -e ';' filename** excludes all lines that start with either a **#** or a **;**.
6. Use **?** to match one or more of the preceding characters.
7. **grep -i text file** finds both *text* and *TEXT* in a file.

8. **grep -A5 '^PATH' filename** shows all lines starting with *PATH* as well as the preceding five lines.
9. **sed -n 9p ~/samplefile** shows line 9 from ~/samplefile.
10. **sed -i 's/user/users/g' ~/samplefile** replaces the word *user* with the word *users* in ~/samplefile.

Chapter 5

1. Typically, the console is the main screen on a Linux server.
2. Press Ctrl-Alt-F2 to switch back from a text-based login prompt to current work on the GUI.
3. **w** or **who** shows all users who currently have a terminal session open to a Linux server.
4. */dev/pts/0* is the device name that is used by the first SSH session that is opened to a server where no GUI is operational.
5. **ssh -v** shows detailed information on what SSH is doing while logging in.
6. **ssh -X** initiates an SSH session with support for graphical applications.
7. *~/.ssh/ssh_config* needs to be edited to modify SSH client settings.
8. **scp /etc/hosts lisa@server2:/tmp** copies the */etc/hosts* file to the directory */tmp* on *server2* using the username *lisa*.
9. *~/.ssh/authorized_keys* stores public keys for remote users who want to log in using key-based authentication.
10. **ssh-keygen** generates an SSH public/private key pair.

Chapter 6

1. 0 is the UID of user root.
2. **sudo** is defined in */etc/sudoers*.
3. Use **visudo** to modify a **sudo** configuration.
4. */etc/default/useradd* and */etc/login.defs* can be used to define settings that will be used when creating users.
5. None, because groups are created in */etc/group*.
6. Making a user a member of the wheel group grants the user access to all admin commands through **sudo**.

7. Use **vigr** to modify the `/etc/group` file manually.
8. **passwd** and **chage** can be used to change user password information.
9. `/etc/shadow` stores user passwords.
10. `/etc/group` stores group accounts.

Chapter 7

1. **chown :groupname filename** or **chown .groupname filename** sets the group owner to a file.
2. **find / -user username** finds all files owned by a specific user.
3. **chmod -R 770 /data** applies read, write, and execute permissions to all files in `/data` for the user and group owners while setting no permissions for others.
4. In relative permission mode, use **chmod +x file** to add the execute permission to a file that you want to make executable.
5. Using **chmod g+s /directory** ensures that group ownership of all new files created in a directory is set to the group owner of that directory.
6. **chmod +t /directory** ensures that users can only delete files of which they are the owner or files that are in a directory of which they are the owner.
7. **setfacl -m g:sales:r *** adds an ACL that grants members of the group sales read permissions to all existing files in the current directory.
8. Using the commands **setfacl -R -m g:sales:rx /dir** and **setfacl -m d:g:sales:rx /dir** ensures that members of sales get read permissions to all files in the current directory and all of its subdirectories, as well as all files created in this directory in the future.
9. Make sure that the last digit in the umask is a 7.
10. **chattr +i myfile** ensures that nobody can delete myfile by accident.

Chapter 8

1. 213.214.215.96 is the network address in 213.214.215.99/29.
2. **ip link show** shows link status and not the IP address.
3. NetworkManager manages network configuration in RHEL 8.
4. `/etc/hostname` contains the hostname in RHEL 8.

5. **hostnamectl** enables you to set the hostname easily.
6. Run **nmcli con up** after manually changing the contents of the `/etc/sysconfig/ifcfg` files.
7. Change `/etc/hosts` to enable hostname resolution for a specific IP address.
8. **ip route show** shows the current routing configuration.
9. **systemctl status NetworkManager** verifies the service's current status.
10. **nmcli con mod "static" ipv4.addresses "10.0.0.20/24" 10.0.0.100** changes the current IP address and default gateway on your network connection.

Chapter 9

1. **createrepo** enables you to make a directory containing a collection of RPM packages a repository.
2. The line `[some-label] name=some-name baseurl=http://server.example.com/repo` needs to be in the repository file.
3. **yum repolist** verifies that a repository is available.
4. **yum provides */useradd** enables you to search the RPM package containing the file `useradd`.
5. Using **yum group list** followed by **yum group info "Security Tools"** shows the name and contents of the yum group that contains security tools.
6. **yum module enable php:5.1** ensures that all PHP-related packages are going to be installed using the older version 5.1, without actually installing anything yet.
7. **rpm -q --scripts packagename** enables you to ensure that a downloaded RPM package does not contain dangerous script code.
8. **rpm -qd name-of.rpm** shows all documentation in an RPM package.
9. **rpm -qf /path/to/file** shows which RPM package a file comes from.
10. **repoquery** enables you to query software from the repository.

Chapter 10

1. **jobs** gives an overview of all current shell jobs.
2. Press `Ctrl-Z` and type **bg** to stop the current shell job to continue running it in the background.

3. Press **Ctrl-C** to cancel the current shell job.
4. Use process management tools such as **ps** and **kill** to cancel the job.
5. **ps fax** shows parent-child relationships between processes.
6. Use **renice -nn -p PID**, where **nn** is a value between -1 and -20. Notice that you need to be root in order to increase process priority.
7. **killall dd** stops all running processes.
8. **pkill mycommand** stops **mycommand**.
9. **k** is used to kill a process.
10. The **tuned** service must be running to select a performance profile.

Chapter 11

1. A unit is a thing that is started by **systemd**. There are different types of units, such as services, mounts, sockets, and many more.
2. Use **systemctl list-units** to show all service units that are currently loaded.
3. You create a want for a service by using the **systemctl enable** command.
4. Set the **SYSTEMD_EDITOR** variable in **/etc/profile** to change the default editor for **systemctl**.
5. **/etc/systemd/system/** contains custom **systemd** unit files.
6. Include **Requires** to ensure that a unit file will automatically load another unit file.
7. **systemctl show httpd** shows available configuration options for the **httpd.service** unit.
8. **systemctl list-dependencies --reverse** shows all dependencies for a specific unit.
9. **systemctl status** output indicating that a unit is dead is nothing serious; it simply means the service is currently not running.
10. Using **systemctl edit** on the unit that you want to modify creates a **systemd** override file.

Chapter 12

1. A cron job that needs to be executed once every 2 weeks is configured as a specific cron file in **/etc/cron.d**, or tied to a user account using **crontab -e -u username**.

2. **0 14 1,15 * *** specifies the execution time in a cron job that needs to be executed twice every month, on the 1st and the 15th of the month at 2 p.m.
3. ***/2 * * * *** specifies the cron execution time for a job to run every 2 minutes every day.
4. **0 0 19 9 4** specifies a job that needs to be executed on September 19 and every Thursday in September.
5. **sun, 0,** and **7** all specify that a cron job needs to be executed on Sunday.
6. **crontab -e -u lisa** enables you to schedule a cron job for user lisa.
7. Create the file `/etc/cron.deny` and make sure that it includes username boris.
8. Specify the job in `/etc/anacrontab` and make sure that the anacron service is operational.
9. The `atd` service must be running to schedule at jobs; use **systemctl status atd** to verify.
10. Use **atq** to find out whether any current at jobs are scheduled for execution.

Chapter 13

1. `/etc/rsyslog.conf` is used to configure rsyslogd.
2. `/var/log/secure` contains messages related to authentication.
3. Log files are rotated away by default after five weeks (one week for the current file, and four weeks for old files).
4. **logger -p user.notice "some text"** logs a message from the command line to the user facility, using the notice priority.
5. Create a file in `/etc/rsyslog.d`. The name does not really matter. Give it the following contents: `*.=info /var/log/messages.info`.
6. You can configure the journal to grow beyond its default size restrictions in `/etc/systemd/journald.conf`.
7. **journalctl -f** shows new messages in the journal scroll by in real time.
8. **journalctl _PID=1 --since 9:00:00 --until 15:00:00** shows all journald messages that have been written for PID 1 between 9:00 a.m. and 3:00 p.m.
9. On a system where a persistent journal has been configured, **journalctl -b** shows journald messages since the last reboot.

10. Making the journald journal persistent requires the following four commands, in order: **mkdir /var/log/journal**; **chown root:systemd-journal /var/log/journal**; **chmod 2755 /var/log/journal**; **killall -USR1 systemd-journald**.

Chapter 14

1. **gdisk** is used to create GUID partitions.
2. **fdisk** is used to create MBR partitions.
3. XFS is the default file system on RHEL 8.
4. **/etc/fstab** is used to automatically mount partitions while booting.
5. The **noauto** mount option is used to specify a file system should not be mounted automatically while booting.
6. **mkswap** enables you to format a partition that has type 82 with the appropriate file system.
7. **mount -a** enables you to test, without actually rebooting, whether automatic mounting of the partitions while booting is going to work.
8. Ext2 is created if you use the **mkfs** command without specifying a file system.
9. Use either **mkfs.ext4** or **mkfs -t ext4** to format an Ext4 partition.
10. Use **blkid** to find UUIDs for all devices on your computer.

Chapter 15

1. The 8e00 partition type is used on a GUID partition that needs to be used in LVM.
2. **vgcreate vgggroup -s 4MiB /dev/sdb3** creates the specified volume group.
3. **pvs** shows a short summary of the physical volumes on your system as well as the volume group to which these belong.
4. Just type **vgextend vgggroup /dev/sdd**. You do not have to do anything on the disk device itself.
5. Use **lvcreate -L 6M -n lvvol1 vggname**. Notice that this works only if you have created the volume group with a 2-MiB physical extent size.
6. **lvextend -L +100M /dev/vgname/lvvol1** adds 100 MB to the logical volume lvvol1.
7. `UUID=xxx /stratis1. Xfs. defaults,x-systemd.requires=stratisd.service 0 0`

8. `stratis pool create mypool /dev/sdd` creates a Stratis pool that is based on the block device `/dev/sdd`.
9. `vdo create --name=vdo1 --device=/dev/sde --vdoLogicalSize=1T` creates a VDO device with a virtual size of 1 TiB.
10. Use the discard option `-k` with `mkfs.xfs` to create an XFS file system on top of a VDO thin-provisioned volume.

Chapter 16

1. `uname -r` shows the current version of the kernel on a computer.
2. Current version information about your RHEL 8 installation is found in `/etc/redhat-release`.
3. `lsmod` shows a list of currently loaded kernel modules.
4. `modinfo modulename` displays kernel module parameters.
5. `modprobe -r` unloads a kernel module.
6. Use `lsmod` to find out which other kernel modules currently need this kernel module and unload these kernel modules first. Note that this will not always work, especially if the considered hardware currently is in use.
7. Use `modinfo` to find which kernel module parameters are supported.
8. Create a file in `/etc/modprobe.d` and include the parameters using an `options` statement.
9. Include `options cdrom debug=1` in the file that will automatically load the `cdrom` module.
10. `yum upgrade kernel` installs a new version of the kernel.

Chapter 17

1. A unit is a thing that is started by `systemd`. There are different types of units, such as services, mounts, sockets, and many more.
2. Use `systemctl mask` to make sure that a target is no longer eligible for automatic start on system boot.
3. Modify `/etc/default/grub` to apply common changes to GRUB 2.
4. `systemctl --type=service` shows all service units that are currently loaded.

5. Create a want for a service by using **systemctl enable** on that service.
6. **systemctl isolate rescue.target** switches the current operational target to the rescue target.
7. There are two types of targets: targets that can run independently and targets that cannot. Check the target unit file to find out more about this and ensure the target is isolatable (which means it can run independently).
8. **systemctl list-dependencies --reverse** shows which other units have dependencies to a systemd service.
9. Apply changes to GRUB 2 in `/etc/default/grub`.
10. Run **grub2-mkconfig > /boot/grub2/grub.cfg** after applying changes to the GRUB 2 configuration.

Chapter 18

1. Press **e** to enter the GRUB boot menu editor mode.
2. An error in `/etc/fstab` prevents the **fsck** command on that file system from finishing successfully.
3. Pass **systemd.unit=rescue.target** to the kernel to enter a mode that provides access to most of the machine’s functionality.
4. Start from a rescue system.
5. **systemctl list-units** shows which units are available in a specific troubleshooting environment.
6. Pass **rd.break** to the GRUB 2 boot loader to start troubleshooting a lost root password.
7. **load_policy -i** loads the SELinux policy.
8. **chcon -t shadow_t /etc/shadow** ensures that the SELinux labels are set correctly.
9. **grub2-mkconfig -o /boot/grub2/grub.cfg** saves changes applied to the GRUB 2 boot loader.
10. **systemd.unit=emergency.target** enters the most minimal troubleshooting mode.

Chapter 19

1. The script will be interpreted by the same shell as the parent shell.
2. **test -z \$VAR** or **[-z \$VAR]** can be used to check whether a variable VAR has no value.
3. Use **\$#** to count the number of arguments that have been used.
4. Use **\$@** to refer to all arguments that have been used when starting the script.
5. Use **read SOMEVAR** to process user input in a script.
6. **[-f filename] || echo file does not exist** determines whether the file exists and, if not, executes the specified command.
7. **[-e filename]** can be used to determine whether an item is a file or a directory.
8. A **for** statement is typically used to evaluate a range of items.
9. You do not; it is a part of the **if** statement that is closed with **fi**.
10. Using **;;** after the last command closes the specific item.

Chapter 20

1. Use **ssh-agent** and **ssh-add** to cache the passphrase that is set on your private key.
2. Use **AllowUsers lisa** to disallow root login and only allow user lisa to log in to your server.
3. Specify the **Port** line twice to configure your SSH server to listen on two different ports.
4. The main SSH configuration file is **/etc/ssh/sshd_config**.
5. The passphrase will be stored in a protected area in memory.
6. **/etc/ssh/ssh_config** contains SSH client settings for all users.
7. The **MaxSessions** parameter that manages this feature is already set to 10 as a default, so you don't need to change anything.
8. **semanage port -a -t ssh_port_t -p tcp 2022** configures SELinux to allow SSH to bind to port 2022.
9. **firewall-cmd --add-port 2022/tcp --permanent; firewall-cmd --reload** configure the firewall on the SSH server to allow incoming connections to port 2022.
10. Try **UseDNS**. This option, which is active by default, uses DNS to get the name of the target host for verification purposes.

Chapter 21

1. The Basic Web Server group contains useful Apache packages.
2. **systemctl enable --now httpd** starts the httpd service automatically when booting.
3. `/etc/httpd/conf.d` is the default location where RPMs can drop plug-in configuration files that should be considered by the Apache server.
4. **curl** enables you to test a web server from a server that does not offer a graphical interface.
5. `/etc/httpd/conf/httpd.conf` is the default Apache configuration file.
6. `/var/www/html` is used as the default Apache document root.
7. The Apache process looks for `index.html`.
8. Use either **systemctl status httpd** or **ps aux | grep http** to check whether the Apache web server is currently running.
9. `/etc/httpd/conf.d` is the preferred location for storing virtual host configuration files.
10. The `ServerRoot` is set by default to `/etc/httpd`.

Chapter 22

1. **setenforce 0** puts SELinux in permissive mode temporarily.
2. **getenforce -a** or **semanage boolean -l** provides a list of all available Booleans.
3. Use the **sepolicy manpage** command to see the service-specific SELinux man page.
4. Install `setroubleshoot-server` to get easy-to-read SELinux log messages in the audit log.
5. Use **semanage fcontext -a -t httpd_sys_content_t "/web(/.*)?"** followed by **restorecon** to apply the `httpd_sys_content_t` context type to the directory `/web`.
6. Never!
7. Change `/etc/sysconfig/selinux` to completely disable SELinux.
8. SELinux logs all of its messages in `/var/log/audit/audit.log`.
9. **man -k _selinux | grep ftp** shows which context types are available for the ftp service.
10. Use **setenforce 0** to temporarily switch SELinux to permissive mode and try again.

Chapter 23

1. **firewalld** should be running before you try to create a firewall configuration with **firewall-config**.
2. **firewall-cmd --add-port=2345/udp** adds UDP port 2345 to the firewall configuration in the default zone.
3. **firewall-cmd --list-all-zones** lists all firewall configuration in all zones.
4. **firewall-cmd --remove-service=vnc-server** removes the vnc-server service from the current firewall configuration.
5. **--reload** activates a new configuration added with the **--permanent** option.
6. **--list-all** enables you to verify that a new configuration has been added to the current zone and is now active.
7. **firewall-cmd --add-interface=eno1 --zone=public** adds the interface eno1 to the public zone.
8. The new interface will be added to the default zone.
9. **firewall-cmd --permanent --add-source=192.168.0.0/24** adds the source IP address 192.168.0.0/24 to the default zone.
10. **firewall-cmd --get-services** lists all services that are currently available in firewalld.

Chapter 24

1. The **showmount** command needs the **mountd** and **rpc.bind** services to be opened in the firewall as well.
2. **showmount -e server1** shows available NFS mounts on server1. Note that the **showmount** command does not get through a firewall.
3. **mount [-t nfs] server1:/share /somewhere** mounts an NFS share that is available on server1:/share.
4. **smbclient -L** shows SMB mounts on a specific server.
5. **cifs-utils** must be installed on an SMB client before you can make an SMB mount.
6. **mount -t cifs -o guest //server1/data /mnt** mounts the Samba share data on server1 with guest access on the local directory /mnt.
7. Use a credentials file that contains at least the username and password that need to be used. Specify all mount options and use **creds=/somewhere/credentials.file** in the mount options column.

8. `auto.master` is the main automount configuration file.
9. `autofs` implements automount.
10. None. You'll have to open ports on the server, not on the client.

Chapter 25

1. `date -s 16:24` sets the system time to 4:24 p.m.
2. `hwclock --systohc` sets the hardware time to the current system time.
3. `date -d '@nnnnnnnn'` shows epoch time as human-readable time.
4. `hwclock --hctosys` synchronizes the system clock with the hardware time.
5. `chronyd` is used to manage NTP time on RHEL 8.
6. `timedatectl set-ntp 1` enables you to use NTP time on your server.
7. `/etc/chrony.conf` contains the list of NTP servers to be used.
8. Either `timedatectl list-timezones` or `tzselect` can be used to list time zones.
9. `timedatectl set-timezone ZONE` is used to set the current time zone.
10. `timedatectl set-time TIME` is used to set the system time.

This page intentionally left blank

Glossary

\$PATH A variable that contains a list of directories that are searched for executable files when a user enters a command.

- The current directory. Its value can be requested using the **pwd** command.

A

absolute filename A filename that is complete and starts with the name of the root directory, including all directories up to the current file or directory.

access control list (ACL) In Linux permissions, a system that makes it possible to grant permissions to more than one user and more than one group. Access control lists also allow administrators to set default permissions for specific directories.

anacron A service that ensures that vital cron jobs can be executed when the server is down at the moment that the job normally should be executed. Can be considered an extension to cron.

AND A logical construction that can be used in scripts. In an AND construction, the second command is executed only after successful execution of the first command.

Application profile A collection of packages that may be used to install a specific version of software, according to a specific installation profile.

Application Stream A specific version of a yum module that can be installed as such.

archiving A system that ensures that data can be properly backed up.

at A service that can be used to schedule future jobs for one-time execution.

attribute A property that can be set to a file or directory and that will be enforced no matter which user with access permission accesses the file. For instance, a file that has the immutable (**i**) attribute set cannot be deleted, not even by the root user. However, the root user does have the capability to change the attribute, which would allow the root user to delete the file anyway.

audit log The main log file in `/var/log/audit/audit.log`, which by default contains all messages that are logged by the `auditd` service.

auditd A service that runs by default on Red Hat Enterprise Linux and can be configured to log very detailed information about what is happening on RHEL. Auditing is complementary to system logging and can be used for compliancy reasons. On RHEL, the auditing system takes care of logging SELinux-related messages, which makes it a relatively important system.

autofs A service that takes care of automatically mounting file systems at the moment that a specific directory is accessed. This service is very useful to ensure the automatic mounting of home directories for users in a centralized user management system, as can be implemented by the LDAP service.

automount The process that is started by the `autofs` service. *See* `autofs` for more details.

B

background process A process that is running on a system without actively occupying a console. Processes can be started in the background by adding a `&` after the command that starts the process. *See also* foreground process.

backup A copy of important data, which can be restored if at any point in time the original data gets lost.

Bash The default shell that is used on Red Hat Enterprise Linux.

Basic Input Output System (BIOS) The first software that is started when a computer starts on older IBM-compatible computers. Settings in the BIOS can be changed by using the BIOS setup program. *See also* Unified Extensible Firmware Interface (UEFI).

binary A numbering scheme that is based on bit values that can be on or off. Binary numbers are 0 and 1. Because binary numbers are difficult to use, decimal, hexadecimal, or octal numbers often are used.

BIOS *See* Basic Input Output System.

boot loader Program that is started as the very first thing while starting a computer and that takes care of loading the operating system kernel and `initramfs`.

Btrfs A general-purpose Linux file system that is expected to become the default file system on Red Hat Enterprise Linux in a future release.

C

cache In memory management, the area of memory where recently used files are stored. Cache is an important mechanism to speed up reads on servers.

capability A specific task that can be performed on Linux. User root has access to all capabilities; normal users have access to limited sets of capabilities only.

CentOS A Linux distribution that uses all Red Hat packages but has removed the Red Hat logo from all these packages to make it possible to distribute the software for free. CentOS is the best option for practicing for the RHCSA exam if you do not have access to RHEL.

certificate In PKI cryptography, contains the public key of the issuer of the certificate. This public key is signed with the certificate of a certificate authority, which guarantees its reliability.

certificate authority (CA) A commonly known organization that can be used to guarantee the reliability of PKI certificates. The certificate authority provides a certificate that can be used to sign public key certificates. Instead of using commonly known organizations, self-signed certificates can be used for internal purposes as well.

chrony The service that offers time synchronization services in Red Hat Enterprise Linux.

chroot An environment where a part of the file system is presented as if it were the root of the file system. Chroot is used as a security feature that hides part of the operating system that is not required by specific services.

CIFS *See* Common Internet File System.

cloud A computing platform that allows for flexible usage of hosted computing resources.

Common Internet File System (CIFS) The standardized version of the Microsoft Server Message Block (SMB) protocol, which is used to provide access to shared printers, files, and directories in a way that is compatible with Windows servers and clients. CIFS has become the de facto standard for file sharing in IT.

compression A technology that is used to reduce the size of files by analyzing redundant patterns and storing them more efficiently.

conditional loop In shell scripting, a set of commands that is executed only if a specific condition has been met.

connection (in network card configuration) A set of network configuration parameters that is associated to a network interface.

connection (in network communication) A session between two parties that has been initialized and will exist until the moment that the connection is tiered down.

console In Linux, the primary terminal where a user works. It is also a specific device with the name `/dev/console`.

context In SELinux, a label that is used to define the security attributes of users, processes, ports, and directories. These contexts are used in the SELinux policy to define security rules.

context switch When the CPU switches from executing one task to executing another task.

context type In SELinux, a label that identifies the SELinux properties of users, processes, ports, and processes.

Coordinated Universal Time (UTC) A time standard that is globally the same, no matter which specific time zone a user is in. UTC roughly corresponds to Greenwich Mean Time (GMT).

credentials file A file that can be used to mount CIFS file systems automatically from the `/etc/fstab` file. The credentials file is stored in a secure place, like the home directory of user root, and contains the username and password that are used to mount the remote file system.

cron A service that takes care of starting services repeatedly at specific times.

cryptography A technique used to protect data, often by converting information to an unreadable state, where keys are used to decipher the scrambled data. Cryptography is used not only to protect files while in transit but also to secure the authentication procedure.

D

deduplication A storage technology that analyzes data to be stored on disk and takes out duplicate patterns to allow for more efficient storage. Used in VDO. *See also* Virtual Data Optimizer (VDO).

default route The route that is used by default to forward IP packets to that have a destination on an external network.

dependency Generally, a situation where one item needs another item. Dependencies occur on multiple levels in Linux. In RPM package management, a dependency is a software package that needs to be present for another package to be installed. In systemd, a dependency is a systemd unit that must be loaded before another unit can be loaded.

dependency hell Situation where for package installation, other packages are needed, which by themselves could require dependencies as well. The problem of dependency hell has been fixed by the introduction of repository-based systems.

destination In rsyslog, specifies where log messages should be sent by the logging system. Destinations are often files, but can also be input modules, output modules, users, or hosts.

device A peripheral that is attached to a computer to perform a specific task.

device file A file that is created in the /dev directory and that is used to represent and interact with a device.

device mapper A service that is used by the Linux kernel to communicate with storage devices. Device mapper is used by LVM, multipath, and other devices, but not by regular hard disks. Device files that are created by device mapper can be found in the /dev/mapper directory.

directory A folder in the file system that can be used to store files in an organized manner.

disabled mode The SELinux mode in which SELinux is completely deactivated.

distribution A Linux version that comes with its own installation program or which is ready for usage. Because Linux is a collection of different tools and other components, the Linux distribution gathers these tools and other components, may or may not enhance them, and distributes them so that users do not have to gather all the different components for themselves.

dmesg Utility that can be used to read the kernel ring buffer, which contains log messages that were generated by the Linux kernel.

dracut A utility that is used to generate the initramfs, an essential part of the Linux operating system that contains drivers and other vital files required to start a Linux system.

dynamic route A network route that is managed by an automatic routing protocol.

E

enforcing mode The SELinux mode where SELinux is fully operational and applies all restrictions that have been configured for a specific system.

environment The collection of settings that users or processes are using to do their work.

epoch time In Linux, the number of seconds that have passed since epoch (corresponds to midnight on January 1, 1970). Some utilities write epoch time instead of real clock time.

escaping In a shell environment, using special syntax to ensure that specific characters are not interpreted by the shell. Escaping may be necessary to show specific characters onscreen or to ensure that regular expression metacharacters are not interpreted by the bash shell first.

export In NFS, refers to a directory that is shared on an NFS server to allow access to other servers.

Ext2, 3, and 4 Three different versions of the Ext file system. Up to RHEL 6, Ext4 was the default file system. It is now considered inadequate for modern storage needs, which is why Ext4 in RHEL 7 has been replaced by XFS as the default file system.

extended partition A solution to create more than four partitions on an MBR disk. On MBR disks, a maximum of four partitions can be stored in the partition table. To make it possible to go beyond that number, one of the four partitions can be created as an extended partition. Within an extended partition, logical partitions can be created, which will perform just like regular partitions, allowing system administrators to create more partitions.

external command A command that exists as a file on disk.

F

facility In rsyslogd, the source where log information comes from. A strictly limited number of facilities have been defined in rsyslogd.

Fedora The free and open source Linux distribution that is sponsored by Red Hat. In Fedora, new features are provided and tested. Some of these features will be included in later releases of Red Hat Enterprise Linux.

FHS *See* Filesystem Hierarchy Standard.

file descriptor A pointer that is used by a Linux process to refer to files that are in use by the process.

file system A logical structure that is created on a storage device. In a Linux file system, inodes are used for file system administration, and the actual data is written to blocks. *See also* inode.

Filesystem Hierarchy Standard (FHS) A standard that defines which Linux directories should be used for which purpose. Read man 7 hier for a specification of the FSH.

firewall A solution that can be used to filter packets on a network. Firewalls are used to ensure that only authorized traffic can reach a system. A firewall can be offered through the Linux kernel Netfilter functionality but often is also offered as an appliance on the network.

firewalld The modern service (replacing iptables) that is used in RHEL 7 and RHEL 8 to implement firewalling based on the Linux kernel firewalling framework.

folder Also referred to as a directory, a structure in the file system used to organize files that belong together.

foreground process Linux processes that are started by users can be started in the foreground or in the background. If a process has been started as a foreground process, no other processes can be started in the same terminal until it finishes or is moved to the background. *See also* background process.

fstab A configuration file that is used on Linux to mount file systems automatically when the system starts.

G

GECOS A field in the `/etc/passwd` file that can be used to store personal data about a user on the Linux operating system. GECOS originally stood for General Electric Comprehensive Operating Supervisor.

global unique ID (GUID) An identification number that consists of parts that ensure that it is globally unique.

GPT *See* GUID Partition Table.

group A collection of items. In user management, groups are used to assign permissions to multiple users simultaneously. In Linux, every user is a member of at least one group.

group owner The group that has been set as the owner of a file or a directory. On Linux, every file and directory has a user owner and a group owner. Group ownership is set when files are created, and unless configured otherwise, it is set to the primary group of the user who creates the file.

GRUB *See* GRUB 2.

GRUB 2 The boot loader that is installed on most systems that need to start Linux. GRUB 2 provides a boot prompt from which different kernel boot options can be entered, which is useful if you need to troubleshoot the boot procedure.

GUID *See* global unique ID.

GUID Partition Table (GPT) A modern solution to store partitions on a hard disk, as opposed to the older MBR partition table. In GUID partitions, a total of 128 partitions can be created, and no difference exists between primary, extended, and logical partitions anymore.

gzip One of the most common utilities that is used for compression and decompression of files on Linux.

H

hard link A name associated with an inode. Inodes are used to store Linux files. An inode contains the complete administration of the file, including the blocks in which the file is stored. A file that does not have at least one hard link is considered a deleted file. To increase file accessibility, more than one hard link can be created for an inode.

hardware time The time that is provided by computer hardware, typically the BIOS clock. When a Linux system boots, it sets the software time based on the hardware time. Because hardware time often is inaccurate, most Linux systems use the Network Time Protocol (NTP) to synchronize the system time with a reliable time source.

hexadecimal A 16-based numbering system that is based on groups of 4 bytes. Hexadecimal numbers start with the range 0 through 9, followed by A through F. Because hexadecimal is much more efficient in computer technology, hexadecimal numbers are often used. In IPv6, IP addresses are written as hexadecimal numbers.

hypervisor A piece of computer software, firmware, or hardware that creates and runs virtual machines. In Linux, KVM is used as the common hypervisor software.

I

inheritance In permission management, refers to the situation where new files that are created in a directory inherit the permission settings from the parent directory.

init The first process that is started once the Linux kernel and initramfs have been loaded. From the init process, all other processes are started. As of RHEL 7, the init process has been replaced by systemd.

initramfs The initial RAM file system. Contains drivers and other files that are needed in the first stages of booting a Linux system. On Red Hat Enterprise Linux, the initramfs is generated during installation and can be manually re-created using the **dracut** utility.

inode Contains the complete administration of a file. Every Linux file has an inode, and the inode contains all properties of the file but not the filename.

input module In rsyslog, a module that allows rsyslog to receive log messages from specific sources.

interface In Linux networking, the set of configuration parameters that can be activated for a specific device. Several interface configurations can exist for a device, but only one interface can be active at a time for a device.

internal command A command that is a part of the shell and does not exist as a file on disk.

Internet Protocol (IP) The primary communications protocol that is used by computers for communication. The Internet Protocol exists in two versions (IPv4 and IPv6). Apart from node addressing, it defines routing, which enables nodes to contact one another.

IP *See* Internet Protocol.

iptables An older solution to create firewall rules on the Linux operating system. Interfaces with the Netfilter Linux kernel firewalling functionality and was the default solution to create software firewalls on earlier versions of RHEL. As of RHEL 7, iptables has been replaced by firewalld.

IPv4 Version 4 of the Internet protocol. It was developed in the 1970s and introduced in 1981. It allows a theoretical maximum of about 4 billion nodes to be addressed by using a 32-bit address space. It is still the most important IP version in use.

IPv6 Version 6 of the Internet protocol. It was developed in the 1990s to address the shortage in IPv6 addresses. It uses a 128-bit address space that allows for addressing 3,4e38 nodes and thus is considered a virtually unlimited address space.

iteration In shell scripting, one time of many that a conditional loop has been processed until the desired result has been reached.

J

job In a Linux shell, a task running in the current terminal. Jobs can be started in the foreground and in the background. Every job is also visible as a process.

journalctl The part of systemd that takes care of logging messages.

journald The part of systemd that takes care of logging information about events that have been happening. The introduction of journald ensures that information about all services can be logged, regardless of how the service itself is configured to deal with information that is to be logged.

K

kernel The central component of the operating system. It manages I/O requests from software and translates them into data processing instructions for the hardware in the computer.

kernel ring buffer A part of memory where messages that are generated by the kernel are stored. The **dmesg** command enables you to read the contents of the kernel ring buffer.

kernel space The part of memory that is reserved for running privileged instructions. Kernel space is typically accessible by the operating system kernel, kernel extensions, and most device drivers. Applications normally run in user space, which ensures that a faulty application cannot crash the computer system.

Kernel-based Virtual Machine (KVM) The Linux kernel module that acts as a hypervisor and makes it possible to run virtual machines directly on top of the Linux kernel.

key-based login In SSH, login that uses public/private keys to proof the identity of the user who wants to log in. Key-based login is generally considered more secure than password-based login.

kill A command that can be used to send a signal to a Linux process. Many signals are defined (see `man 7 signal`), but only a few are commonly used, including `SIGTERM` and `SIGKILL` that both are used to stop processes.

KVM *See* Kernel-based Virtual Machine.

L

label A name that can be assigned to a file system. Using labels can be a good idea, because once a label is assigned, it will never be changed, which guarantees that the file system can still be mounted, even if other parameters such as the device name have changed. However, UUIDs are considered safer than labels because the chance of having a duplicate label by accident is much higher than the chance of having a duplicate UUID. *See also* universally unique ID (UUID).

line anchor In regular expressions, a character that refers to a specific position in a line.

Linux A UNIX-like operating system that consists of a kernel that was originally developed by Linus Torvalds (hence the name Linux). A current Linux operating system consists of a kernel and lots of open source tools that provide a complete operating system. Linux is packaged in the form of a distribution. Currently, Red Hat Enterprise Linux is among the most widely used Linux distributions.

log rotation A service that ensures that log files cannot grow too big. Log files are monitored according to specific parameters, such as a maximum age or size.

Once this parameter is reached, the log file will be closed and a new log file will be opened. Old log files are kept for a limited period and will be removed, often after only a couple of weeks.

logical extent The building block that is used in LVM to create logical volumes. It normally has a size of a few megabytes that corresponds to the size of the physical extents that are used.

logical partition A partition that is created in an extended partition. *See also* extended partition.

logical volume In LVM, the entity on which a file system is created. Logical volumes are often used on RHEL because they offer important advantages, such as the option to dynamically resize the logical volume and the file system that it hosts.

Logical Volume Manager (LVM) The software that makes it possible to work with logical volumes.

login shell The shell that is opened directly after a user has logged in.

LVM *See* Logical Volume Manager.

M

masquerading A solution that enables a private IP address range that is not directly accessible from outside networks to be accessed by using one public IP address that is exposed on a router.

Master Boot Record (MBR) On a BIOS system, the first 512 bytes on the primary hard disk. It contains a boot loader and a partition table that give access to the different partitions on the hard disk of that computer.

MBR *See* Master Boot Record.

module A piece of snap-in code. Modules are used by several systems on Linux, such as the kernel, GRUB 2, rsyslog, and more. Via modules, Linux components can be extended easily, and adding functionality does not require a total rewrite of the software.

module (in YUM) A collection of software packages that can be managed as one entity and can contain different versions of a software solution.

mount A connection that is made between a device and a directory. To access files on specific storage devices, the storage device needs to be mounted on a directory. This sets up the specified directory as the access point to files on the storage device. Mounts are typically organized by the systems administrator and are not visible to end users.

multiplier In regular expressions, a character that indicates that multiple of the previous character are referred to.

N

netfilter The part of the Linux kernel that implements firewalling.

netmask *See* subnet mask.

Network Address Translation (NAT) *See* masquerading.

Network File System (NFS) A common UNIX solution to export physical file systems to other hosts on the network. The other hosts can mount the exported NFS directory in their local file system.

network time Time that is provided on the network.

Network Time Protocol (NTP) A standard that is used to provide reliable time to servers in a network. NTP on RHEL 8 is implemented by the chronyd service.

NFS *See* Network File System.

nftables The service that manages kernel firewalling. It is a replacement of the older iptables service.

nice A method to change the priority of Linux processes. A negative nice value will make the process more aggressive, giving it a higher priority (which is expressed by a lower priority number); a positive nice value will make a process less eager so that it gives priority to other processes.

NTP *See* Network Time Protocol.

O

octal A numbering scheme that uses the numbers 0 through 7 only. Used when working with Linux permissions using the **umask** or the **chmod** commands.

OR A logical operation where the second command is executed only if the first command is not able to execute.

output module In rsyslog, a module that is used to send log messages to a specific destination. Output modules make rsyslogd flexible and allow for the usage of log destinations that are not native to rsyslog.

ownership In file system permissions, the basis of the effective permissions that a user has. Every file has a user owner and a group owner assigned to it.

P

package A bundle that is used to distribute software. A package typically contains a compressed archive of files and metadata that includes instructions on how to install those files.

package group A group of packages that can be installed as such using the **yum groups install** command.

package group (in yum) A group of software packages that can be installed with a single command.

pager A program that can be used to browse page by page through a text file. The **less** utility provides one of the most common Linux pagers.

parent shell The environment from which a shell script or program is started. Processes or child scripts will inherit settings from the parent shell.

partition A subdivision of a hard disk on which a file system can be created to mount it into the directory structure.

passphrase Basically a password, but is supposed to be longer and more secure than a password.

password A token that is used in authentication. The password is a secret word that can be set by individual users and will be stored in an encrypted way.

path The complete reference to the location of a file.

permissions Attributes that can be set on files or directories to allow users or groups access to these files or directories.

permissive mode The SELinux mode where nothing is blocked but everything is logged in the audit log. This mode is typically used for troubleshooting SELinux issues.

physical extent The physical building block that is used when creating LVM physical volumes. Typically, the size is multiple megabytes.

physical volume The foundation building block of an LVM configuration. The physical volume typically corresponds to a partition or a complete disk device.

PID *See* process identification number.

pipe A structure that can be used to forward the output of one command to be used as input for another command.

policy *See* SELinux policy.

port A number that is used by a process to offer access to the process through a network connection.

port forwarding A firewalling technique where traffic that is coming in on a specific port is forwarded to another port that may be on the same host or on a different host.

Portable Operating System Interface (POSIX) A standard that was created to maintain compatibility between operating systems. The standard mainly applies to UNIX and guarantees that different flavors of Linux and UNIX are compatible with one another.

portmapper A remote procedure call service that needs to run on systems that provide RPC services. Portmapper uses dynamic ports that do not correspond to specific TCP or UDP ports; the service will pick a UDP or TCP port that will be used as long as the process is active. When restarted, chances are that different ports are used. They need to be mapped to fixed UDP and TCP ports in order to make it possible to open the firewall for these ports. Portmapper is still used by components of the NFS service.

POSIX *See* Portable Operating System Interface.

primary group The group that is listed in the group membership field for a user in `/etc/passwd`. Every Linux user is a member of a primary group. Apart from that, users can be made a member of secondary groups as well.

primary partition In MBR, one of a maximum of four partitions that can be created in the Master Boot Record. *See* also extended partition.

priority (in process handling) Specifies the importance of a process. Process priority is expressed with a number (which can be modified using **nice**). Processes with a lower priority number are serviced before processes with a higher priority number.

priority (in rsyslog) Used to specify the severity of a logged event. Based on the severity, specific actions can be taken.

private key In public/private key encryption, the key that is used to generate encrypted data.

privileged user *See* root.

proc A kernel interface that provides access to kernel information and kernel tunables. This interface is available through the `/proc` file system.

process A task that is running on a Linux machine. Roughly, a process corresponds to a program, although one program can start multiple processes.

process identification number (PID) A unique number that is used to identify a process running on a Linux system.

profile In tuned, a collection of performance settings that can easily be applied.

protocol A set of rules that is used in computing, such as in computer networking to establish communications between two computers.

Pseudo Root File System In the current NFS release NFSv4, a solution where multiple shares are exported by an NFS server. Instead of mounting each individual share, the NFS client mounts the root file system on the NFS server, which gives access to all shares the client is entitled to.

pseudo root mount In NFSv4, a mount of the root directory system or another high-level directory that hasn't been specifically exported by the NFS server, but which gives access to all exported file systems the client has access to.

public key In cryptography, the key that is typically sent by a server to a client so that the client can send back encrypted data.

PV *See* physical volume.

Q

queue In process management, where processes wait before they can be executed.

R

real-time clock The hardware clock that is installed on the computer motherboard.

reboot The procedure of stopping the computer and starting it again.

Red Hat Customer Portal The platform that Red Hat offers to provide patches for customers that have an active subscription. To provide these patches and updates, Red Hat Network provides the repositories that are needed for this purpose.

Red Hat Enterprise Linux (RHEL) The name of the software that Red Hat sells subscriptions for. It is available in a server edition and a desktop edition.

Red Hat Package Manager The name for the package format that is used on RHEL for software packages and for the Package Management software. RPM has become the standard for package management on many other Linux distributions as well.

reference clock A clock that is used as a time source in an NTP time configuration. Typically, a reference clock is a highly reliable clock on the Internet, but it can be an internal clock on the computer's motherboard as well.

regular expression A search pattern that allows users to search text patterns in a flexible way. Not to be confused with shell metacharacters.

relative filename A filename that is relative to a directory that is not the root directory.

Remote Procedure Calls (RPC) A method for interprocess communication that allows a program to execute code in another address space. Remote Procedure Calls is an old protocol and as such is still used in the Network File System.

repository An installation source that contains installable packages and an index that contains information about the installable packages so that the installation program yum can compare the version of packages currently installed with the version of packages available in the repository.

resident memory Memory pages that are in use by a program.

resolver The DNS client part that contains a list of DNS servers to contact to resolve DNS queries.

RHEL *See* Red Hat Enterprise Linux.

rich rules Rules in firewalld that allow the usage of a more complicated syntax so that more complex rules can be defined.

root The privileged user account that is used for system administration tasks. User root has access to all capabilities, which means that permissions do not apply to user root and the root user account is virtually unlimited.

root directory The starting point of the file system hierarchy, noted as /.

RPC *See* Remote Procedure Calls.

RPM *See* Red Hat Package Manager.

RTC *See* real-time clock.

rsyslogd The generic daemon that logs messages.

S

Samba The name for the Linux service that implements the SMB protocol.

SAN *See* storage-area network.

scheduler The part of the Linux kernel that monitors the queue of runnable processes and allocates CPU time to these processes.

Scientific Linux A Linux distribution that is based on the Red Hat packages from which the Red Hat logo has been removed. A very good choice for people that are looking for a freely available alternative to Red Hat Enterprise Linux.

secondary group A group that a user is a member of but which membership is not defined in the /etc/passwd file. When creating new files, the secondary group will not automatically become the owner of those files.

Secure Shell (SSH) A solution that allows users to open a shell on a remote server where security is implemented by using public/private key cryptography.

Secure Sockets Layer (SSL) *See* Transport Layer Security (TLS).

SELinux A Linux kernel security module that provides a mechanism for supporting access control security policies.

SELinux Policy The collection of rules that is used to define SELinux security.

Server Message Block (SMB) An application-level protocol that is used to provide shared access to files, printers, and serial ports, which on Linux is implemented in the Samba server.

services (in firewalld) A configuration of firewall settings that is used to allow access to specific processes.

services (in systemd) Processes that need to be started to provide specific functionality.

share A directory to which remote access is configured using a remote file system protocol such as NFS or CIFS.

shebang Used in a script to indicate which shell should be used for executing the code in the shell script. If no shebang is used, the script code will be interpreted by the parent shell, which may lead to errors in some cases. A shebang starts with a #, which is followed by a ! and the complete pathname of the shell, such as #!/bin/bash.

shell The environment from which commands can be executed. Bash is the default shell on Linux, but other shells exist as well.

shell metacharacters Characters such as *, ?, and [a-z] that allow users to refer to characters in filenames in a flexible way.

signal An instruction that can be sent to a process. Common signals exist, such as SIGTERM and SIGKILL, but the Linux kernel allows a total of 32 different signals to be used. To send a signal to a process, use the **kill** command.

SMB *See* Server Message Block.

snapshot A “photo” of the actual state of a file system.

software time *See* system time.

source context In SELinux, the context of the processes or users that initiate an action. A context in SELinux is a label that identifies allowed operations. Everything in an SELinux environment has a context.

SSH *See* Secure Shell.

standard error (STDERR) The default location where a program sends error messages.

standard input (STDIN) The default location where a program gets its input.

standard output (STDOUT) The default location where a program sends its regular output.

static route A route that is defined manually by a network administrator.

STDERR *See* standard error.

STDIN *See* standard input.

STDOUT *See* standard output.

storage-area network (SAN) A solution where disk devices are shared at a block level over the network. As such, they can be used in the same way as local disk devices on a Linux system. iSCSI and Fibre Channel are the common SAN protocols.

Stratis The new volume managing file system in RHEL 8.

stratum In time synchronization, used to indicate the distance between a server and an authoritative Internet time source.

subnet mask A logical subdivision of an IP network.

subshell A shell that is started from another shell. Typically, by running a shell script a subshell is started.

symbolic link A special type of file that contains a reference to another file or directory in the form of an absolute or relative path.

sysfs The kernel interface that is mounted on the /sys directory and which is used to provide access to parameters that can be used for managing hardware settings.

system time The time that is maintained by the operating system. When a Linux system boots, system time is set to the current hardware time, and while the operating system is running, it is often synchronized using the Network Time Protocol (NTP).

systemd The service manager on RHEL 8. systemd is the very first process that starts after the kernel has loaded, and it takes care of starting all other processes and services on a Linux system.

T

tainted kernel A kernel in which unsupported kernel modules have been loaded.

tar The Tape Archiver; the default Linux utility that is used to create and extract backups.

target In systemd, a collection of unit files that can be managed together.

target context The SELinux context that is set to a target object, such as a port, file, or directory.

terminal Originally, the screen that was used by a user to type commands on. On modern Linux systems, pseudo terminals can be used as a replacement. A pseudo terminal offers a shell window from which users enter the commands that need to be executed.

thin allocation In storage, an approach that enables the system to present more storage to the storage user than what is really available by using smart technologies to store data, like deduplication.

thread A thread can be used as a subdivision of a process. Many processes are single threaded, which means that process is basically one entity that needs to be serviced. On a multicore or multi-CPU computer system, working with multithreaded processes makes sense. That way, the different cores can be used to handle the different threads, which allows a process to benefit from multicore or multithreaded environments.

time stamp An identifier that can be used on files, database records, and other types of data to identify when the last modification has been applied. Many services rely on time stamps. To ensure that time stamped-based systems work properly, time synchronization needs to be configured.

time synchronization A system that ensures that multiple servers are using the exact same time. To accomplish time synchronization, it is common to use an external time server, as defined in the Network Time Protocol (NTP).

timer A systemd unit type that can be used as an alternative to cron jobs and run units at a specific time.

TLS See Transport Layer Security.

Transport Layer Security (TLS) A cryptographic protocol that is created to ensure secured communications over a computer network. In TLS, public and private keys are used, and certificates authenticate the counterparty. TLS was formerly known as SSL.

TTY A program that provides a virtual terminal on Linux. Every terminal still has a TTY name, which is either tty1-6 for virtual TTYs or /dev/pts/0-nn for pseudo terminals.

tuned A service on RHEL that enables administrators to easily apply performance settings by using profiles.

U

udev A service that works with the Linux kernel to initialize hardware.

UEFI *See* Unified Extensible Firmware Interface.

umask An octal value that defines the default permissions as a shell property.

umount The command that is used to decouple a file system from the directory on which it is mounted.

Unified Extensible Firmware Interface (UEFI) A replacement of the Basic Input Output System used on older IBM-compatible computers as the first program that runs when the computer is started. UEFI is the layer between the operating system and the computer firmware.

unit In systemd, refers to an item that is managed by systemd. Different types of units exist, including service, path, mount, and target units.

universally unique ID (UUID) An identification number that consists of a long random hexadecimal number and which is globally unique.

unprivileged user A regular non-root user account to which access restrictions apply, as applied by permissions.

user An entity that is used on Linux to provide access to specific system resources. Users can be used to represent people, but many services also have a dedicated user account, which allows the service to run with the specific permissions that are needed for that service.

user space The area of memory that is accessible by application software that has been started with non-root privileges.

UTC *See* Coordinated Universal Time.

UUID *See* universally unique ID.

V

value The data that is assigned to a specific property, variable, or record.

variable A label that contains a specific value that can be changed dynamically. In scripting, variables are frequently used to allow the script to be flexible.

VDO See Virtual Data Optimizer.

VFAT The Linux kernel driver that is used to access FAT-based file systems. FAT is a commonly used file system in Windows environments. The Linux VFAT driver allows usage of this file system.

VG See volume group.

Virtual Data Optimizer (VDO) A new advanced storage solution in RHEL 8 that compresses data by using deduplication and can therefore be used for thin allocation of storage volumes.

virtual host In the Apache web server, a collection of configuration settings that is used to address a web server. What makes it a virtual host is that one installation of the Apache web server can be configured with multiple virtual hosts, which allows administrators to run multiple websites on one Apache server.

virtual memory The total amount of addressable memory. Virtual memory is called *virtual* memory because it does not refer to memory that really exists. Its only purpose is to make sure that Linux programs can set an address pointer that is unique and not in use by other programs.

volume group The abstraction layer that in Logical Volume Manager is used to represent all available storage presented by physical volumes from which logical volumes can be created.

W

want An indication for a systemd unit file that it is supposed to be started from a specific systemd target.

wildcard The * character, which in a shell environment refers to an unlimited number of any characters.

X

XFS A high-performance 64-bit file system that was created in 1993 by SGI and which in RHEL 8 is used as the default file system.

Y

Yellowdog Update, Modified The full name for YUM, the meta package handler that on RHEL 8 is used to install packages from yum repositories.

Yum *See* Yellowdog Update, Modified.

Z

zombie A process that has lost contact with its parent and for that reason cannot be managed using regular tools.

zone In firewalld, a collection of one or more network interfaces that specific firewalld rules are associated with.

Index

SYMBOLS

- & (ampersand) command, 237–238
- \ (backslash) escape character, 433
- ~/.bash_profile, 43
- ~/.bashrc, 43
- /boot directory, 56–57
- /dev directory, 56
- . (dot), 553, 91
- /etc directory, 56
- /etc/anacrontab, 278–279
- /etc/bashrc, 43
- /etc/crontab file, managing, 276–278
- /etc/default/grub, 395
- /etc/default/useradd, 132–133
- /etc/dracut.conf, 412–413
- /etc/fstab, 332–335
 - NFS share mounting, 520–521
 - Samba share mounting, 521
- /etc/group, 136–137
- /etc/gshadow, 137
- /etc/hosts, 191–192
- /etc/httpd, 464–465
- /etc/httpd/conf/httpd.conf, 460–463
- /etc/issue, 43
- /etc/login.defs, 133
- /etc/logrotate.conf, 298
- /etc/motd, 43
- /etc/passwd, 89
 - fields, 128–129
 - modifying, 131
- /etc/profile, 43
- /etc/shadow
 - fields, 129–130
 - modifying, 131
- /etc/sysconfig/selinux, 476
- /etc/systemd/journald.conf, 303
- help option, 44–45
- /home directory, 56, 58
- && (logical AND), 553, 433
- || (logical OR), 564, 433
- /media directory, 56
- /mnt directory, 56
- /opt directory, 56
- \$PATH, 553, 33
- /proc directory, 56, 372
- / (root) directory, 56
- /root directory, 56
- /run directory, 56
- #! (shebang), 569, 426
- /srv directory, 56
- /sys directory, 56
- /tmp directory, 57
- /usr directory, 57–58
- /usr/share/doc documentation files, 49
- /var directory, 57–58
- /var/log, list of log files, 290
- /var/log/messages, 290–292
- * (wildcard), 573

A

- absolute filenames
 - defined, 553
 - overview, 63–64
- accessing repositories, 8

accounts. *See* group accounts; user accounts

ACLs (access control lists), 156–160
 changing/viewing settings, 157–159
 default ACLs, 159–160
 defined, 553
 file system preparation, 157
 support for, 156–157

advanced permissions, 152–156

aliases, 33

alternative port configuration, 447

ampersand (&) command, 237–238

anacron

configuring, 278–279

defined, 553

AND, 553, 433

Apache servers

configuration files, 460–465

content creation, 463–464

software requirements, 460

verifying availability, 464

virtual hosts, 465–467

Application profile, 553

Application Stream, 553, 220

apropos command, 47

archiving

defined, 554

files

extracting tar files, 73

with tar, 72

arguments

defined, 33

in shell scripts, 428–429

asterisk (*) wildcard, 573

at command, 281

defined, 554

task scheduling, 281

atq command, 281

atrm command, 281

attributes

defined, 554

user-extended, 162–163

audit log

defined, 554

SELinux messages in, 488–490

auditd, 554

authentication

key-based, 114–115, 451–452

to Samba shares, 518

autofs, 521–523

configuring, 522

defined, 554

defining mounts, 522

overview, 521

wildcards in, 523

automount, 521–523

configuring, 522

defined, 554

defining mounts, 522

overview, 521

wildcards in, 523

awk command, 93–94

B

background processes

defined, 554

running, 236–237

types of, 239

backslash (\) escape character, 433

backups

of ACLs, 157

defined, 554

Bash. *See also* shell scripting

commands

command-line completion, 39

executing, 33

help with, 44–49

history, 37–38

I/O redirection, 34–35

pipes, 36

syntax, 32–33

defined, 554, 32

escaping, 90

bash -x command, 438

bash-completion, 39

Basic Input Output System (BIOS), 554

- batch command, 281
- bg command, 237, 238
- binary, 554
- binary notation
 - for IP addresses, 172
- BIOS (Basic Input Output System), 554
- blkid command, 331–332
- Boolean settings for SELinux, 487–488
- boot loaders, 554
- boot process
 - GRUB 2, 394–397
 - changing configuration, 394–395
 - configuration files, 395–396
 - modifying default options, 396–397
 - steps in, 404–405
 - systemd targets
 - enabling services, 390–391
 - isolating, 391–394
 - setting default, 394
 - target units, 389–390
 - tasks of, 388–389
 - types of, 388
 - viewing, 391–393
 - wants, 390
- troubleshooting
 - accessing boot prompt, 406–407
 - file system issues, 414–415
 - overview, 405
 - recovering virtual machine access, 416–418
 - re-creating initramfs image, 414
 - reinstalling GRUB 2, 413–414
 - rescue disk usage, 408–413
 - resetting root password, 416
 - starting troubleshooting targets, 407–408
- boot prompt, accessing, 406–407
- broadcast addresses, 171
- Btrfs, 554
- bunzip2 command, 74
- bzip2 command, 74

C

- cache, 555
- cache tier, 356
- capabilities, 555
- case statement, 436–437
- cat command, 84–85
- CentOS
 - defined, 555
 - installing, 10–24
 - obtaining, 7
- Cert Guide environment setup, 9–10
- certificate authority (CA), 555
- certificates, 555
- chage command, 133–134
- characters
 - counting, 88–89
 - escaping, 433
- chattr command, 163
- chcon command, 481
- chgrp command, 148
- chmod command, 150–152, 155
- chown command
 - group ownership, 148
 - user ownership, 147–148
- chrony
 - defined, 555
 - time service client configuration, 537–538
- chroot, 555
- chroot /mnt/sysimage command, 409–410, 414
- chvt command, 105
- CIFS (Common Internet File System)
 - defined, 555
 - history of, 516–517
 - mounting through fstab, 521
 - server configuration, 518–519
 - shares
 - authentication, 518
 - discovering, 517–518, 520
 - mounting, 518, 520

- cloud, 555
- command mode (vim), 40
- command-line completion, 39
- commands
 - command-line completion, 39
 - executing, 33
 - help with, 44–49
 - /usr/share/doc documentation files, 49
 - help option, 44–45
 - info command, 48–49
 - man command, 45–48
 - mandb command, 47–48
 - history, 37–38
 - I/O redirection, 34–35
 - LVM management, 353
 - pipes, 36
 - syntax, 32–33
 - vim editor, 40–41
- Common Internet File System. *See* CIFS (Common Internet File System)
- compression
 - defined, 555
 - overview, 73–74
- conditional loops, 432–437
 - case statement, 436–437
 - defined, 555
 - for loops, 429, 434–435
 - if.then.else construction, 432–434
 - until loops, 435–436
 - while loops, 435–436
- configuration files
 - for Apache servers, 460–465
 - GRUB 2, 395–396
 - ifcfg, 189
 - modifying, 131
 - rsyslogd, 293, 297
 - in shell environment, 43
 - for user management, 132–133
- connections
 - defined, 555
 - devices versus, 180
 - fixed and dynamic IP addresses for, 189
 - SSH server options, 450
 - viewing properties, 182–185
 - viewing status, 181
- consoles
 - defined, 555, 102
 - local consoles
 - logging in, 102–103
 - pseudo terminals, 106
 - switching terminals, 103–105
 - remote access
 - with Secure Shell, 108–110, 112
 - transferring files, 111–112
- context labels, 479–481
- context switches, 556
- context types, 481–485
 - defined, 556
 - finding via man pages, 484–485
 - restoring defaults, 485–486
 - setting, 481–484
- contexts, 556
- Coordinated Universal Time (UTC), 556, 530
- copying files
 - cp command, 65–66
 - scp command, 112
- counting words/lines/characters, 88–89
- cp command, 65–66
- createrepo command, 208
- credentials files, 556
- cron
 - anacron service, 278–279
 - components of, 274
 - configuration file management, 276–278
 - crond management, 274–275
 - defined, 556
 - exercise, 280
 - security, 280
 - timing, 276
- crond, managing, 274–275
- crontab command, 277

cryptography, 556
cut command, 84, 87

D

daemons, 236
date command, 532
daylight saving time, 530
debugging shell scripts, 438
deduplication, 556
default ACLs, 159–160
default boot options, modifying, 396–397
default file contexts, restoring, 485–486
default ownership, 148–149
default permissions, 160–161
default routes, 556
default targets, setting, 394
default zones in firewalld, 501
deleting files, 66
dependencies
 defined, 556
 of units, managing, 263–264
dependency hell, 556, 225
destinations
 defined, 556
 in rsyslogd, 294–296
device files, 557, 35
device mappers, 557
devices
 connections versus, 180
 defined, 557
 mounting, 57–61, 331–332
 naming, 351–352
df -hT command, 61
df -Th command, 59
dictionary attacks, 446
direct write logging, 288
directories. *See also* files
 defined, 557
 in FSH, 56–57
 home, 132
 listing, 64–65
 managing, 62–63
 mounting, 57–58
 ownership
 changing groups, 148
 changing users, 147–148
 default ownership, 148–149
 viewing, 146–147
 permissions
 ACLs, 156–160
 advanced, 152–156
 basic, 149–152
 default, 160–161
 user-extended attributes, 162–163
disabled mode, 557
disabling
 root login, 446
 SELinux, 478–479
disk device types, 315
disk labels, mounting with, 331–332
distributions, 557
dmesg
 analyzing kernel activity, 371–372
 defined, 557
dnf command, yum command and, 200
DNS name resolution, 192–193
documentation files, 49
dot (.), 553, 91
downloading GPG keys, 206–207
dracut command
 defined, 557
 re-creating initramfs image, 411–413, 414
drivers, 371, 374–375
 checking availability, 379–380
dynamic IP addresses, 174, 189
dynamic routes, 557

E

echo command, 43
editing files with vim, 39–41
elevated permissions
 running tasks with, 125
 PolicyKit, 126

- su command, 125
 - sudo command, 126
- emergency reset, 108
- emergency.target, 388
- enabling
 - modules, 223
 - SELinux, 475–476
 - services, 390–391
- enforcing mode, 557
- entitlements, 201
- env command, 42
- environments, 557
- EPEL (Extra Packages for Enterprise Linux) repositories, 201
- epoch time, 557, 532
- escaping
 - characters, 433
 - defined, 557
 - regular expressions, 90
- exam preparation, 541–544
 - nondisclosure agreement (NDA), 544
 - registering for exam, 541–542
 - SELinux coverage, 492
 - theoretical pre-assessment exam, 545–547
 - tips for taking exam, 542–544
 - verifying readiness, 541
 - what to bring to exam, 542
- examples
 - /etc/crontab file, 277
 - /etc/default/grub, 395
 - /etc/default/useradd, 132–133
 - /etc/dracut.conf, 412–413
 - /etc/fstab, 333
 - /etc/group, 136
 - /etc/hosts, 191–192
 - /etc/httpd, 464
 - /etc/httpd/conf/httpd.conf, 461–463
 - /etc/logrotate.conf, 298
 - /etc/passwd, 89, 128
 - /etc/shadow, 129
 - /etc/sysconfig/selinux, 476

- /etc/systemd/journald.conf, 303
- /var/log/messages, 291
- anacron configuration, 279
- analyzing kernel activity, 372
- arguments in shell scripts, 428
- basic script example, 426
- blkid command, 332
- case statement, 437
- chage command, 134
- CIFS share discovery, 517
- connection properties, 182–185
- connection status, 181
- context label viewing, 480
- counting with wc command, 89
- cron job in /etc/cron.d, 278
- current environment display, 42
- debugging shell scripts, 438
- df -hT command, 61
- file ownership display, 147
- file system formatting with XFS, 326–327
- filtering with cut command, 87
- findmnt command, 59–60
- firewall configuration, 504
- firewalld services, 502
- for loops, 434
- ftp service file, 503
- GPG key downloads, 206–207
- group ACL changes, 158
- hostname configuration, 190–191
- hostnamectl status command, 373
- id command, 124
- ifcfg configuration files, 189
- if.then.else construction, 433
- ip addr show command, 176
- ip link show command, 177
- ip route show command, 178–179
- journalctl -o verbose command, 301–302
- kernel thread viewing, 239–240
- line anchors, 90
- link properties, 70

- listing kernel threads, 370
- log file viewing from journald, 299–300
- log file viewing with systemctl status command, 289
- logging in with Secure Shell, 110
- lsblk command, 349
- lsmod command, 376
- lspci command, 379–380
- LVM device naming, 352
- modinfo command, 377–378
- monitoring time synchronization status, 534–535
- mount command, 58–59
- mount units, 258
- multi-user.target, 389
- newgrp command, 149
- permissions checking, 158
- physical volume verification, 348
- ps aux command, 88
- ps -ef command, 240–241
- ps fax command, 241–242
- pvcreeate command, 345
- pvdisplay command, 348
- read command in shell scripts, 430
- regular expression necessity, 89–90
- repository availability verification, 209
- repository file, 204–205
- RULES section of rsyslog.conf, 294
- screen command, 115–116
- sealert command, 490–491
- searching man pages, 46
- SELinux messages in audit.log, 489
- semanage fcontext usage, 483
- service units, 257
- sestatus command, 477
- socket units, 259
- ss -lt command, 179
- systemctl list-dependencies command, 263–264
- systemctl show command, 265
- systemctl status command, 262
- systemctl status crond -l command, 275
- systemd mount units for VDO, 360–361
- systemd targets, 391–393
- target units, 260
- time zone settings, 536–537
- timedatectl command, 533
- tune2fs command, 327–328
- udevadm monitor command, 375
- unit types in systemd, 256
- until loops, 436
- uptime command, 247
- vgdisplay command, 350
- volume groups resizing verification, 354
- web server availability verification, 464
- while loops, 435
- yum groups info command, 218
- yum groups list command, 217
- yum history command, 219
- yum info nmap command, 211–212
- yum install nmap command, 212–213
- yum list command, 214
- yum list kernel command, 215
- yum module info command, 223
- yum module list command, 221–222
- yum module list perl command, 223
- yum search command, 210
- yum update command, 215–216
- execute permissions, 149–152
- executing commands, 33
- exercises
 - ACLs (access control lists), 160
 - advanced permissions, 156
 - Apache virtual host configuration, 466–467
 - at task scheduling, 281
 - automount configuration for NFS, 522
 - bash-completion, 39
 - basic permission management, 152
 - connection parameter changes with nmcli, 186
 - context label setting, 483–484

- cron task scheduling, 280
- directory management, 62–63
- file management, 67–68
- file system creation, 327
- firewall management with
 - firewall-cmd, 503–505
- GPT partition creation with gdisk, 321–324
- GPT partition creation with parted, 325
- grep options, 93
- group account management, 137
- GRUB 2 modifications, 397
- head command, 86
- history (in Bash), 38
- if.then.else construction, 434
- info command, 49
- input in shell scripts, 431
- internal and external commands, 34
- I/O redirection and pipes, 36–37
- isolating targets, 394
- journalctl command, 300–301
- kernel module management, 378
- kernel modules with parameters, 380
- less command, 85
- live log file monitoring, 292
- load average management, 248
- local time management, 534
- logging in with Secure Shell, 109–110
- logical partition creation, 319–320
- logical volume resizing, 355–356
- man -k command, 48
- MBR partition creation with fdisk, 316–318
- module management, 224–225
- mounting partitions through /etc/fstab, 335
- mounts overview, 60
- network configuration validation, 177
- network connections management with nmcli, 186
- network settings verification, 179–180
- NFS share mounting, 516
- NFS share offering, 515
- package management with yum, 219
- physical volume creation, 346–347
- positional parameters, 429
- preserving systemd journal, 303
- process management from command line, 245
- pseudo terminals, 106
- remote access with public/private keys, 114–115
- remote access with Secure Shell, 112
- repository creation, 208–209
- rescuing system, 410–411
- restorecon command, 486
- RPM queries, 228–229
- rsyslog.conf rule changes, 297
- Samba server configuration, 519
- SELinux Boolean settings, 488
- SELinux mode manipulation, 478
- SELinux-specific man pages, 485
- shell environment management, 44
- shell job management, 238
- simple shell script creation, 427
- SMB share discovery and mounting, 520
- SSH security option configuration, 449
- Stratis volume management, 358–359
- swap partition creation, 330
- switching terminals, 104
- switching user accounts, 127
- symbolic links and hard links, 71
- tail command, 86
- tar, 75
- time service client configuration, 538
- troubleshooting targets, 408
- tuned, 249
- unit configuration changes, 266
- unit management with systemctl, 261
- user account creation, 135
- VDO storage management, 361
- vim practice, 41

- volume group and logical volume
 - creation, 352
- web server setup, 463
- wildcard automount configuration, 523
- exit command, 427
- exports, 557
- Ext2, 557
- Ext3, 557
- Ext4
 - defined, 557
 - property management, 327–329
- extended partitions
 - defined, 557
 - on MBR, 319–320
- external commands, 558, 33
- Extra Packages for Enterprise Linux (EPEL) repositories, 201
- extracting tar files, 73

F

- facilities
 - defined, 558
 - in rsyslogd, 294–296
- fdisk command, 315–320
- Fedora
 - defined, 558
 - obtaining, 8
- fg command, 237, 238
- FHS (Filesystem Hierarchy Standard)
 - defined, 558
 - mounts, 57–61
 - overview, 56–57
- file command, 73
- file descriptors, 558
- file systems. *See also* storage
 - ACL preparation, 157
 - automount, 521–523
 - configuring, 522
 - defining mounts, 522
 - overview, 521
 - wildcards in, 523

CIFS

- authentication, 518
- discovering shares, 517–518, 520
- history of, 516–517
- mounting shares, 518, 520
- mounting through fstab, 521
- server configuration, 518–519
- creating, 325–327
- defined, 558
- mounting, 330–335
 - automating with /etc/fstab, 332–335
 - with device name/UUID/disk label, 331–332
 - manually, 331

NFS

- automount configuration, 522
- mounting shares, 515–516
- mounting through fstab, 520–521
- security, 514
- versions, 514–515
- property management, 327–329
- troubleshooting, 414–415
- types of, 326

files. *See also* directories; text files

- absolute and relative file names, 63–64
- archiving
 - extracting tar files, 73
 - with tar, 72
- compressing, 73–74
- copying
 - cp command, 65–66
 - scp command, 112
- deleting, 66
- editing with vim, 39–41
- links
 - creating, 69–70
 - hard links, 68–69, 71
 - removing, 70–71
 - symbolic links, 69, 71
- listing, 64–65
- moving, 66
- ownership

- changing groups, 148
- changing users, 147–148
- default ownership, 148–149
- viewing, 146–147
- permissions
 - ACLs, 156–160
 - advanced, 152–156
 - basic, 149–152
 - default, 160–161
 - user-extended attributes, 162–163
- synchronizing, 112
- transferring securely, 111–112
- wildcards, 61–62
- Filesystem Hierarchy Standard. *See* FHS (Filesystem Hierarchy Standard)
- filtering text files
 - with cut command, 87
 - with head command, 86
 - with tail command, 86
- find command, 147
- finding
 - context types, 484–485
 - packages, 210–211
- findmnt command, 59–60
- firewall-cmd command, 503–506
- firewalld, 500–503
 - defined, 558
 - services, 502–503
 - zones, 501
- firewalls
 - defined, 558
 - firewall-cmd command, 503–506
 - firewalld, 500–503
 - services, 502–503
 - zones, 501
 - netfilter, 500
 - viewing configuration, 504
- first lines of text files, viewing, 86
- fixed IP addresses, 174, 189
- folders, 558. *See also* directories

- foreground processes
 - defined, 559
 - running, 236–237
- for loops, 429, 434–435
- fsck command, 414
- fstab, 559, 332–335
 - NFS share mounting, 520–521
 - Samba share mounting, 521

G

- gdisk command, 320–324
- GECOS, 559
- getenforce command, 476
- getfacl command, 157–159
- getsebool command, 487
- GiB, 9
- global unique ID (GUID), 559
- GPG keys, downloading, 206–207
- GPT (GUID Partition Table)
 - creating
 - with gdisk command, 320–324
 - with parted command, 324–325
 - defined, 559
 - overview, 313
- graphical applications with Secure Shell, 111
- graphical.target, 388
- grep command, 242
 - options, 92–93
- group accounts
 - creating, 136–137
 - groupadd command, 137
 - vigr command, 136–137
 - membership checking, 137
 - property management, 137
 - types of, 135–136
- group owners
 - changing, 148
 - defined, 559
- groupadd command, 137

- groupmems command, 137
- groupmod command, 137
- groups
 - defined, 559
 - of packages, 216–218
- groups command, 148
- GRUB 2, 394–397
 - accessing boot prompt, 406–407
 - changing configuration, 394–395
 - configuration files, 395–396
 - defined, 559
 - modifying default options, 396–397
 - reinstalling, 411, 413–414
- grub2-install command, 413–414
- grub2-mkconfig command, 396–397
- GUID (global unique ID), 559
- GUID Partition Table. *See* GPT (GUID Partition Table)
- gunzip command, 74
- gzip command, 559, 74

H

- hard links
 - defined, 560
 - exercise, 71
 - overview, 68–69
- hardware clock, 530
- hardware initialization, 371, 374–375
 - checking driver availability, 379–380
- hardware time, 560
- head command, 84, 86
- help with commands, 44–49
 - /usr/share/doc documentation files, 49
 - help option, 44–45
 - info command, 48–49
 - man command, 45–48
 - mandb command, 47–48
- hexadecimal, 560
- history
 - in Bash, 37–38
 - of CIFS, 516–517
 - of packages, 218–219

- home directories, 132
- hostnamectl set-hostname command, 190
- hostnamectl status command, 190, 373
- hostnames, 190–192
- hosts, 170. *See also* virtual hosts
- httpd package, 460
- hwclock command, 532
- hypervisors, 560

I

- id command, 124
- ifcfg configuration files, 189
- ifconfig command, 175
- if.then.else construction, 432–434
- info command, 48–49
- inheritance, 560
- init, 560
- initramfs
 - defined, 560
 - re-creating from rescue disk, 411–414
- inodes, 560, 68
- input in shell scripts, 430–431
- input mode (vim), 40
- input modules, 560
- insmod command, 378
- installing
 - modules, 224
 - packages, 212–213
 - RHEL
 - advantages of subscription, 6
 - Cert Guide environment, 9–10
 - free alternatives, 7
 - manual installation, 10–24
 - obtaining, 7
 - repository access, 8
 - setup requirements, 9
- interfaces, 560
- internal commands, 561, 33
- Internet Protocol (IP), 561
- I/O redirection, 34–35
- IP (Internet Protocol), 561
- ip addr command, 175

- ip addr show command, 175–176
- ip command, 175–177
- ip link command, 175
- ip link show command, 177
- ip route command, 175
- ip route show command, 178
- iptables command, 561, 500
- IPv4 addresses
 - binary notation, 172
 - defined, 561, 170
 - fixed versus dynamic, 174
 - overview, 170–171
 - subnet masks, 171
 - validating configuration, 175–177
- IPv6 addresses
 - defined, 561, 170
 - overview, 171
- isolating systemd targets, 391–394
- iterations, 561

J

- jobs
 - defined, 561
 - shell jobs. *See* shell jobs
- jobs command, 237–238
- journalctl command, 561, 288, 299–302
- journalctl -o verbose command, 301–302
- journald, 299–303
 - defined, 561, 288
 - journalctl command, 299–302
 - overview, 288–289
 - preserving systemd journal, 302–303

K

- kernel
 - analyzing activity, 371–373
 - defined, 561, 370
 - drivers, 371, 374–375
 - upgrading, 381
- kernel modules, 373, 376–380
 - drivers, checking availability, 379–380
 - listing, 376

- loading/unloading, 378
- parameter management, 380
- viewing information about, 377–378
- kernel ring buffer, 561
- kernel space, 561
- kernel threads
 - defined, 236
 - listing, 370
 - showing, 239–240
- Kernel-based Virtual Machine (KVM), 561
- key-based authentication, 114–115, 451–452
- key-based logins, 562
- kill command, 562, 239, 244–245
- killall command, 245
- KVM (Kernel-based Virtual Machine), 561

L

- labels, 562
- last lines of text files, viewing, 86
- less command, 84–85
- line anchors, 562, 90
- lines, counting, 88–89
- links
 - creating, 69–70
 - hard links, 68–69, 71
 - removing, 70–71
 - symbolic links, 69, 71
- Linux, 562
- listing
 - directories, 64–65
 - files, 64–65
 - firewalld services, 502
 - kernel modules, 376
 - kernel threads, 370
 - modules, 221–223
 - packages, 214–215
- live log file monitoring, 292
- ln command, 69–70
- load average for processes, 247

- loading kernel modules, 378
- local consoles
 - logging in, 102–103
 - pseudo terminals, 106
 - switching terminals, 103–105
- local time
 - defined, 530
 - managing, 534
 - time zone settings, 535–537
- log files
 - journal, 299–303
 - journalctl command, 299–302
 - overview, 288–289
 - preserving systemd journal, 302–303
 - rotating, 297–299
 - rsyslogd, 293–297
 - configuration files, 293, 297
 - facilities/priorities/destinations, 294–296
 - overview, 288–289
 - for SELinux
 - in audit log, 488–490
 - sealert command, 490–491
 - types of, 288
 - viewing
 - with journalctl -o verbose command, 301–302
 - from journal, 299–300
 - live monitoring, 292
 - sample content, 290–292
 - with systemctl status command, 289
 - in /var/log, 290
- log rotation
 - defined, 562
 - overview, 297–299
- logger command, 292
- logging in
 - disabling root login, 446
 - with key-based authentication, 114–115
 - to local consoles, 102–103
 - to remote servers, 108–110, 112
- logical AND (&&), 553, 433
- logical extent, 562
- logical OR (||), 564, 433
- logical partitions
 - defined, 562
 - on MBR, 319–320
- Logical Volume Manager. *See* LVM (Logical Volume Manager)
- logical volumes
 - creating, 344–345, 351–352
 - defined, 562
 - resizing, 353–356
- login shells, 562
- loopback interfaces, 176
- loops. *See* conditional loops
- ls command, 64–65, 70
- ls -l command, 146–147
- lsblk command, 348–349
- lscpu command, 247
- lsmod command, 376
- lspci command, 379–380
- lvcreate command, 353
- lvdisplay command, 353
- lvextend command, 354
- LVM (Logical Volume Manager)
 - advantages of, 343–344
 - architecture, 342–343
 - commands, 353
 - defined, 562
 - device naming, 351–352
 - logical volumes
 - creating, 344–345, 351–352
 - resizing, 353–356
 - partitions, creating, 347
 - physical volumes
 - creating, 346–349
 - verifying, 348
 - volume groups
 - creating, 349–350, 352
 - resizing, 353–354

`lvremove` command, 353
`lvresize` command, 354–355
`lvs` command, 353

M

MAC addresses, 173
`man` command, 45–48
`man -k` command, 46–48
`man` pages
 finding context types, 484–485
 searching, 46
 for `semanage` command, 482–483
`mandb` command, 47–48
masquerading, 563
MBR (Master Boot Record)
 creating with `fdisk`, 315–318
 defined, 563
 extended and logical partitions, 319–320
 overview, 312–313
measurement units for storage, 314
`mkfs` command, 326–327
`modinfo` command, 377–378
`modprobe` command, 378
modular kernel. *See* kernel modules
modules
 defined, 563, 220
 enabling, 223
 information about, 223
 installing, 224
 listing, 221–223
 managing, 221–225
 terminology, 220
`more` command, 85
`mount` command, 58–59, 331, 518
mounting
 file systems, 330–335
 automating with `/etc/fstab`, 332–335
 with device name/UUID/disk label, 331–332
 manually, 331
 NFS shares, 515–516, 520–521
 Samba shares, 518, 520–521

mounts, 57–61
 defined, 563
 defining in automount, 522
 in `systemd`, 258
 for VDO, 360–361
 wildcard configuration, 523
moving files, 66
multiplier, 563
`multi-user.target`, 388, 389
`mv` command, 66

N

naming devices, 351–352
nano editor in `systemd`, 266
NAT (Network Address Translation), 171.
 See also masquerading
NDA (nondisclosure agreement), 544
`netfilter`
 defined, 563
 overview, 500
`netmask`. *See* subnet masks
`netstat` command, 178
Network Address Translation (NAT), 171.
 See also masquerading
Network File System. *See* NFS (Network File System)
network masks. *See* subnet masks
network time, 563
Network Time Protocol (NTP)
 defined, 563
 overview, 531
networking
 configuring, 180–189
 with `nmcli` command, 181–187
 with `nmtui` command, 187–189
 permissions for, 180–181
 DNS name resolution, 192–193
 hostnames, 190–192
 IPv4 addresses, 170–171
 binary notation, 172
 fixed versus dynamic, 174
 subnet masks, 171
 validating configuration, 175–177

- IPv6 addresses, 170, 171
 - MAC addresses, 173
 - port addresses, 173
 - validating, 178–179
 - protocols, 173
 - routing, validating, 178
 - verifying settings, 179–180
- newgrp command, 149
- NFS (Network File System)
 - automount configuration, 522
 - defined, 563
 - mounting shares, 515–516
 - mounting through fstab, 520–521
 - security, 514
 - versions, 514–515
- nftables command, 564, 500
- nice command, 564, 242–244
- nmcli command, 180–187
- nm-connection-editor command, 188
- nmtui command, 180, 187–189
- nodes, 170
- nohup command, 239
- nondisclosure agreement (NDA), 544
- normal user accounts, system accounts
 - versus, 127–130
- NTP (Network Time Protocol)
 - defined, 563
 - overview, 531

O

- octal, 564
- optimization with tuned,
 - 248–249
- options, 32
- OR, 564, 433
- output modules, 564
- ownership
 - changing groups, 148
 - changing users, 147–148
 - default, 148–149
 - defined, 564
 - viewing, 146–147

P

- package groups, 564
- Package Module Streams, 220
- packages
 - defined, 564
 - finding, 210–211
 - groups of, 216–218
 - history of, 218–219
 - information about, 211–212
 - installing, 212–213
 - listing, 214–215
 - managing, 219
 - modules. *See* modules
 - Package Module Streams, 220
 - removing, 213–214
 - repoquery command, 228
 - repositories. *See* repositories
 - rpm command, 225
 - database queries, 226
 - filename parts, 226
 - package file queries, 227–228
 - query usage, 228–229
 - updating, 215–216
- paggers, 564
- parent shells, 565
- parent-child relations for shell jobs, 239
- parted command, 324–325, 347
- partitions
 - advantages of multiple, 312
 - defined, 565
 - disk device types, 315
 - GPT (GUID Partition Table), 313
 - gdisk command, 320–324
 - parted command, 324–325
 - LVM partition creation, 347
 - MBR (Master Boot Record), 312–313
 - creating with fdisk, 315–318
 - extended and logical partitions,
 - 319–320
 - swap files, 330
 - swap partitions, 329–330
 - utilities for, 314–315

partprobe command, 320

passphrases

configuring, 451–452

defined, 565

overview, 114

passwd command, 133–134

passwords

defined, 565

dictionary attacks, 446

resetting root password, 416

user account properties, 133–134

paths, 565

performance optimization with tuned,
248–249

permissions

ACLs, 156–160

changing/viewing settings, 157–159

default ACLs, 159–160

file system preparation, 157

support for, 156–157

advanced, 152–156

basic, 149–152

default, 160–161

defined, 565

elevating, 125

PolicyKit, 126

su command, 125

sudo command, 126

networking configuration changes,
180–181

user-extended attributes, 162–163

permissive mode, 565

physical extent, 565

physical volumes

creating, 346–349

defined, 565

verifying, 348

PID (process identification number), 566

pininfo command, 48–49

pipes

defined, 565

less command, 84

overview, 36

pkill command, 245

PolicyKit, 125–126

pools

creating, 357

defined, 356

managing, 358–359

port addresses, 173

validating, 178–179

port forwarding, 565

port scans, avoiding, 447

Portable Operating System Interface
(POSIX), 566

portmappers, 566

ports

alternative port configuration, 447

changing SELinux labels, 447–448

defined, 565

positional parameters in shell scripts,
428–429

POSIX (Portable Operating System
Interface), 566

preparation for exam, 541–544

nondisclosure agreement (NDA), 544

registering for exam, 541–542

theoretical pre-assessment exam,
545–547

tips for taking exam, 542–544

verifying readiness, 541

what to bring to exam, 542

primary groups

changing, 149

defined, 566, 131

primary partition, 566

priorities

defined, 566

of processes, adjusting, 242–244

in rsyslogd, 294–296

private keys, 566

private network addresses,
170–171

privileged users. *See* root

proc, 566

process identification number (PID), 566

- process management
 - from command line, 245
 - kill command, 244–245
 - nice command, 242–244
 - performance optimization, 248–249
 - processes versus threads, 239–240
 - ps command, 240–242
 - shell jobs
 - managing, 237–239
 - parent-child relations, 239
 - running in foreground/background, 236–237
 - top command, 246–248
 - types of processes, 236
- processes
 - adjusting priority, 242–244
 - defined, 566
 - information about, 240–242
 - load average, 247
 - sending signals to, 244–245
 - states of, 247
 - threads versus, 239–240
 - types of, 236
- profiles
 - defined, 566, 220
 - tuned, 248
- programmatic API in Stratis, 356
- protocols
 - defined, 566
 - for networking, 173
- ps aux command, 88, 239–240
- ps command, 240–242
- Pseudo Root File System, 566
- pseudo rout mount, 567
- pseudo terminals, 106
- public keys, 567
- pvcreeate command, 345, 353
- pvdisplay command, 348, 353
- pvremove command, 353
- pvs command, 353
- pwd command, 63

Q

- querying
 - exercise, 228–229
 - repositories, 228
 - RPM database, 226
 - RPM package files, 227–228
- queues, 567

R

- read command, 430
- read permissions, 149–152
- real-time clock, 567, 530
- reboot
 - defined, 567
 - during exam, 544
 - overview, 106–108
- Red Hat Customer Portal, 567
- Red Hat Enterprise Linux. *See* RHEL (Red Hat Enterprise Linux)
- Red Hat Package Manager, 567
- redirectors, 34–35
- reference clock, 567
- registering
 - for exam, 541–542
 - RHEL, 201
- regular expressions, 89–92
 - awk and sed commands, 93–94
 - defined, 567
 - escaping, 90
 - grep options, 92–93
 - line anchors, 90
 - wildcards, 91–92
- reinstalling GRUB 2, 411, 413–414
- relative filenames
 - defined, 568
 - overview, 63–64
- remote access
 - with key-based authentication, 114–115
 - with Secure Shell, 108–110, 112
 - transferring files, 111–112

- rsync command, 112
- scp command, 112
- sftp command, 112
- Remote Procedure Calls (RPC), 568
- removing
 - links, 70–71
 - packages, 213–214
- renice command, 242–244
- repoquery command, 228
- repositories
 - accessing, 8
 - creating, 208–209
 - defined, 568
 - querying, 228
 - role of, 200–201
 - security, 206–207
 - specifying, 202–205
- repository files
 - example, 204–205
 - options, 205
- requirements for RHEL, 9
- rescue disks, 408–413
 - re-creating initramfs image, 411–413, 414
 - reinstalling GRUB 2, 411, 413–414
 - restoring system access, 409–411
- rescue.target, 388
- resident memory, 568
- resizing
 - logical volumes, 353–356
 - volume groups, 353–354
- resolver, 568
- restorecon command, 482, 485–486
- restoring
 - default file contexts, 485–486
 - system access with rescue disk, 409–411
- RHEL (Red Hat Enterprise Linux)
 - defined, 6–7, 567
 - installing manually, 10–24
 - preparation for installation
 - advantages of subscription, 6
 - Cert Guide environment, 9–10
 - free alternatives, 7
 - obtaining, 7
 - repository access, 8
 - setup requirements, 9
 - registering, 201
- rich rules, 568
- rm command, 66
- rmmmod command, 378
- root
 - defined, 568
 - disabling login, 446
 - overview, 124–125
 - PolicyKit, 126
 - resetting password, 416
 - su command, 125
 - sudo command, 126
- root directory, 568
- rotating log files, 297–299
- routing, validating, 178
- RPC (Remote Procedure Calls), 568
- rpm command, 225
 - database queries, 226
 - downloading GPG keys, 206
 - filename parts, 226
 - package file queries, 227–228
 - query usage, 228–229
- RPM packages. *See* packages
- rpm -q --scripts, 227
- rpm -qa command, 226, 228
- rpm -qc command, 226–227
- rpm -qd command, 226–227
- rpm -qf command, 226–227
- rpm -qi command, 226–227
- rpm -ql command, 226–227
- rpm -qp command, 227–228
- rpm -qR command, 227
- rpm -V command, 228
- rpm -Va command, 228
- rsync command, 112
- rsyslogd, 293–297
 - configuration files, 293, 297
 - defined, 568, 288

facilities/priorities/destinations,
294–296
overview, 288–289
RULES section of rsyslog.conf, 294
running processes, 247

S

Samba

defined, 568
history of CIFS, 516–517
server configuration, 518–519
shares
authentication, 518
discovering, 517–518, 520
mounting, 518, 520
mounting through fstab, 521

SAN (storage-area network), 570

scheduler, 568

scheduling tasks

at, 281

cron

anacron service, 278–279
components of, 274
configuration file management,
276–278
crond management,
274–275
exercise, 280
security, 280
timing, 276

Scientific Linux

defined, 568
obtaining, 8

scp command, 112

screen command, 115–116

scripts. *See* shell scripting

sealert command, 490–491

searching

man pages, 46
text files, 89–92

secondary groups, 569, 131

Secure Shell. *See* SSH (Secure Shell)

Secure Sockets Layer (SSL). *See*
Transport Layer Security (TLS)

security

cron, 280

for NFS, 514

repositories, 206–207

SSH server hardening, 446–449

alternative port configuration, 447

disabling root login, 446

limiting user access, 448–449

SELinux port labels, 447–448

sed command, 93–94, 110

SELinux

Boolean settings, 487–488

changing port labels, 447–448

context labels, 479–481

context types, 481–485

finding via man pages, 484–485

restoring defaults, 485–486

setting, 481–484

defined, 569

disabling, 478–479

elements of, 475

enabling, 475–476

exam preparation for, 492

log files

in audit log, 488–490

sealert command, 490–491

modes, 475–479

SELinux Policy, 569

semanage command, 481–483

semanage port command, 447–448

Server Message Block. *See* SMB (Server
Message Block)

services

defined, 569, 256

enabling, 390–391

in firewalld, 502–503

in systemd, 257–258

session options for SSH server, 450

- sestatus command, 476–477
- set group ID (SGID) permission, 153–155
- set user ID (SUID) permission, 152–153, 155
- setenforce command, 476
- setfacl command, 157–159
- setroubleshoot-server package, 491
- setsebool command, 488
- setup requirements for RHEL, 9
- sftp command, 112
- SGID (set group ID) permission, 153–155
- shares
 - CIFS
 - authentication, 518
 - discovering, 517–518, 520
 - mounting, 518, 520
 - mounting through fstab, 521
 - defined, 569
 - NFS
 - mounting, 515–516
 - mounting through fstab, 520–521
- shebang (#!), 569, 426
- shell environment
 - /etc/motd and /etc/issue, 43
 - configuration files, 43
 - variables, 42–43
- shell jobs
 - defined, 236
 - managing, 237–239
 - parent-child relations, 239
 - running in foreground/background, 236–237
- shell metacharacters, 569
- shell scripting
 - arguments, 428–429
 - conditional loops, 432–437
 - case statement, 436–437
 - for loops, 429, 434–435
 - if.then.else construction, 432–434
 - until loops, 435–436
 - while loops, 435–436
 - debugging scripts, 438
 - elements of, 426–427
 - variables, 430–431
- shells, 569, 32. *See also* Bash; shell scripting
- showmount command, 516
- signals
 - defined, 569
 - sending to processes, 244–245
- sleeping processes, 247
- SMB (Server Message Block)
 - defined, 569
 - history of CIFS, 516–517
 - server configuration, 518–519
- shares
 - authentication, 518
 - discovering, 517–518, 520
 - mounting, 518, 520
 - mounting through fstab, 521
- smbclient command, 517–518
- snapshots
 - defined, 569, 344
 - managing, 358
 - in Stratis, 356
- sockets in systemd, 259
- soft links. *See* symbolic links
- software clock, 530
- software management. *See also* yum
 - command
 - modules
 - enabling, 223
 - information about, 223
 - installing, 224
 - listing, 221–223
 - managing, 221–225
 - terminology, 220
 - packages
 - finding, 210–211
 - groups of, 216–218

- history of, 218–219
- information about, 211–212
- installing, 212–213
- listing, 214–215
- managing, 219
- Package Module Streams, 220
- removing, 213–214
- updating, 215–216
- registering RHEL, 201
- repoquery command, 228
- repositories
 - creating, 208–209
 - querying, 228
 - role of, 200–201
 - security, 206–207
 - specifying, 202–205
- rpm command, 225
 - database queries, 226
 - filename parts, 226
 - package file queries, 227–228
 - query usage, 228–229
- subscription management, 201
- software requirements for Apache servers, 460
- software time. *See* system time
- sort command, 84, 87–88
- sorting, text files with sort command, 87–88
- source context, 569
- special permissions. *See* advanced permissions
- ss -lt command, 178–179
- SSH (Secure Shell)
 - connection options, 450
 - defined, 569, 108
 - graphical applications with, 111
 - key-based authentication, 114–115, 451–452
 - list of configuration options, 451
 - remote access with, 108–110, 112
 - server hardening, 446–449
 - alternative port configuration, 447
 - disabling root login, 446
 - limiting user access, 448–449
 - SELinux port labels, 447–448
 - session options, 450
 - transferring files, 111–112
 - rsync command, 112
 - scp command, 112
 - sftp command, 112
 - ssh command, 108–110, 111
 - ssh-add command, 452
 - ssh-agent command, 452
 - ssh-copy-id command, 114
 - sshd. *See* SSH (Secure Shell)
 - ssh-keygen command, 114
 - SSL (Secure Sockets Layer). *See* Transport Layer Security (TLS)
- star command, 73
- static routes, 570
- STDERR (standard error)
 - defined, 570
 - overview, 34
- STDIN (standard input)
 - defined, 570
 - overview, 34
- STDOUT (standard output)
 - defined, 570
 - overview, 34
- sticky bit, 154–155
- stopped processes, 247
- storage. *See also* file systems
 - LVM
 - advantages of, 343–344
 - architecture, 342–343
 - commands, 353
 - device naming, 351–352
 - logical volume creation, 344–345, 351–352
 - logical volume resizing, 353–356
 - partition creation, 347

- physical volume creation, 346–349
- physical volume verification, 348
- volume group creation, 349–350, 352
- volume group resizing, 353–354
- measurement units, 314
- partitions
 - advantages of multiple, 312
 - disk device types, 315
 - formatting with file system, 326–327
 - GPT (GUID Partition Table), 313, 320–325
 - LVM partition creation, 347
 - MBR (Master Boot Record), 312–313, 315–320
 - swap files, 330
 - swap partitions, 329–330
 - utilities for, 314–315
- Stratis
 - architecture, 356–357
 - features of, 356
 - pool creation, 357
 - pool management, 358–359
- VDO
 - configuring, 360–361
 - overview, 359–360
- storage-area network (SAN), 570
- Stratis
 - architecture, 356–357
 - defined, 570
 - features of, 356
 - pools
 - creating, 357
 - managing, 358–359
- stratis blockdev command, 358
- stratis filesystem command, 358
- stratis pool add-data command, 358
- stratis pool command, 358
- stratum, 570
- streams, 220
- su command, 125
- subnet masks
 - defined, 570
 - overview, 171
- subscription management, 201
- subscription-manager tool, 201
- subshells, 570
- sudo command, 125, 126
- SUID (set user ID) permission, 152–153, 155
- superuser. *See* root
- swap files, 330
- swap partitions, 329–330
- switching
 - terminals, 103–105
 - user accounts, 127
- symbolic links
 - defined, 570
 - exercise, 71
 - overview, 69
- synchronizing
 - files, 112
 - time, 534–535
- sysfs, 570
- system accounts, normal user accounts
 - versus, 127–130
- system logging. *See* log files
- system time, 570, 530
- systemctl command, 256, 261–263
- systemctl disable command, 390
- systemctl edit command, 265
- systemctl enable command, 260, 390
- systemctl get-default command, 394
- systemctl halt command, 107
- systemctl isolate command, 393
- systemctl list-dependencies command, 260, 263–264
- systemctl poweroff command, 107
- systemctl set-default command, 394
- systemctl show command, 265
- systemctl status command, 262, 289, 464

- systemctl status crond -l command, 274–275
- systemctl status -l chronyd command, 534–535
- systemctl status NetworkManager command, 180
- systemd
 - defined, 570
 - preserving journal, 302–303
 - rebooting, 107
 - units
 - configuration changes, 266
 - dependency management, 263–264
 - locations, 256–257
 - managing, 261–263
 - mounts, 258, 360–361
 - option management, 265–266
 - services, 257–258
 - sockets, 259
 - targets, 259–260
 - types of, 256
- systemd targets
 - enabling services, 390–391
 - isolating, 391–394
 - setting default, 394
 - target units, 389–390
 - tasks of, 388–389
 - types of, 388
 - viewing, 391–393
 - wants, 390
- systemd-udevd, 374

T

- tac command, 85
- tail command, 84, 86
- tail -f command, 292
- tainted kernel, 571, 371
- Tape ARchiver (tar) utility. *See* tar
- tar, 71–73
 - compressing files, 73–74
 - creating archives, 72
 - defined, 571

- exercise, 75
- extracting files, 73
- options, 74
- target context, 571
- targets
 - defined, 571
 - starting for troubleshooting, 407–408
- in systemd, 259–260
 - enabling services, 390–391
 - isolating, 391–394
 - setting default, 394
 - target units, 389–390
 - tasks of, 388–389
 - types of, 388
 - viewing, 391–393
 - wants, 390
- task scheduling
 - at, 281
 - cron
 - anacron service, 278–279
 - components of, 274
 - configuration file management, 276–278
 - crond management, 274–275
 - exercise, 280
 - security, 280
 - timing, 276
- terminals
 - defined, 571, 102
 - pseudo, 106
 - screen command, 115–116
 - switching, 103–105
 - virtual, 104–105
- test command, 430–431, 432
- text files
 - counting words/lines/characters with wc command, 88–89
 - filtering
 - with cut command, 87
 - with head command, 86
 - with tail command, 86

- regular expressions, 89–92
 - awk and sed commands, 93–94
 - escaping, 90
 - grep options, 92–93
 - line anchors, 90
 - wildcards, 91–92
 - sorting with sort command, 87–88
 - tools for, 80–84
 - viewing
 - with cat command, 85
 - first/last lines of, 86
 - with less command, 84–85
 - theoretical pre-assessment exam, 545–547
 - thin allocation, 571
 - thin provisioning, 356
 - threads
 - defined, 571, 236
 - processes versus, 239–240
 - time
 - date command, 532
 - epoch time, 532
 - hwclock command, 532
 - local time management, 534
 - monitoring synchronization status, 534–535
 - NTP (Network Time Protocol), 531
 - time service clients, 537–538
 - time zone settings, 535–537
 - timedatectl command, 532, 533
 - types of, 530
 - time service clients, configuring, 537–538
 - time stamps, 571
 - time synchronization, 571
 - time zone settings, 535–537
 - timedatectl command, 532–533
 - timers, 571
 - timing in cron, 276
 - TLS (Transport Layer Security), 571
 - top command, 246–248
 - transferring, files securely, 111–112
 - Transport Layer Security (TLS), 571
 - troubleshooting boot process
 - accessing boot prompt, 406–407
 - file system issues, 414–415
 - overview, 405
 - recovering virtual machine access, 416–418
 - re-creating initramfs image, 414
 - reinstalling GRUB 2, 413–414
 - rescue disk usage, 408–413
 - resetting root password, 416
 - starting troubleshooting targets, 407–408
 - TTY, 571
 - tune2fs command, 327–329
 - tuned
 - defined, 571
 - performance optimization, 248–249
 - profile overview, 248
 - tuned-adm command, 249
 - tzselect command, 536–537
- ## U
- udev, 572
 - udevadm monitor command, 374–375
 - UEFI (Unified Extensible Firmware Interface), 572
 - umask
 - defined, 572
 - setting default permissions, 160–161
 - umount command, 572, 331
 - uname command, 373
 - Unified Extensible Firmware Interface (UEFI), 572
 - uninterruptable sleep processes, 247
 - units
 - configuration changes, 266
 - defined, 572
 - dependency management, 263–264
 - locations, 256–257
 - managing, 261–263
 - mounts, 258, 360–361
 - option management, 265–266
 - services, 257–258

- sockets, 259
 - targets, 259–260, 389–390
 - types of, 256
- universally unique ID (UUID)
 - defined, 572
 - mounting with, 331–332
- unloading kernel modules, 378
- unprivileged users, 572
- until loops, 435–436
- updating
 - mandb database, 47–48
 - packages, 215–216
- upgrading kernel, 381
- uptime command, 247
- user accounts
 - creating, 131–132, 135
 - home directories, 132
 - modifying configuration files, 131
 - useradd command, 131
- id command, 124
- limiting SSH server access, 448–449
- password properties, 133–134
- property management, 132–134
- root, 124–125
 - PolicyKit, 126
 - su command, 125
 - sudo command, 126
- switching, 127
- system versus normal accounts, 127–130
- user environment, creating, 134
- user environments, creating, 134
- user ownership, changing, 147–148
- user space, 572
- useradd command, 131
- userdel command, 131
- user-extended attributes, 162–163
- usermod command, 132
- users, 572
- UTC (Coordinated Universal Time), 556, 530

- UUID (universally unique ID)
 - defined, 572
 - mounting with, 331–332

V

- validating
 - IPv4 address configuration, 175–177
 - port availability, 178–179
 - routing, 178
- values, 572
- variables
 - defined, 572
 - in shell environment, 42–43
 - in shell scripts, 430–431
- VDO (Virtual Data Optimizer)
 - configuring, 360–361
 - defined, 572
 - overview, 359–360
- verifying
 - exam readiness, 541
 - network settings, 179–180
 - physical volumes, 348
 - web server availability, 464
- VFAT, 572
- vgcreate command, 349–350, 353
- vgdisplay command, 350, 353
- vgextend command, 353–354
- vgreduce command, 353–354
- vgremove command, 353
- vgs command, 353, 354
- viewing
 - ACL settings, 157–159
 - connection properties, 182–185
 - connection status, 181
 - context labels, 480
 - firewall configuration, 504
 - hostname configuration, 190–191
 - kernel module information, 377–378
 - log files

- with journalctl -o verbose command, 301–302
- from journal, 299–300
- live monitoring, 292
- sample content, 290–292
- with systemctl status command, 289
- in /var/log, 290
- ownership, 146–147
- systemd targets, 391–393
- text files
 - with cat command, 85
 - first/last lines of, 86
 - with less command, 84–85
- vigr command, 136–137
- vim editor
 - commands, 40–41
 - overview, 39–41
 - in systemd, 266
- vimtutor command, 40
- vipw command, 131
- virtual console tty1, 105
- Virtual Data Optimizer. *See* VDO (Virtual Data Optimizer)
- virtual hosts
 - for Apache servers, 465–467
 - defined, 572
- virtual machines (VMs), recovering
 - access, 416–418
- virtual memory, 572
- virtual terminals, 104–105
- visudo command, 126
- VMs (virtual machines), recovering
 - access, 416–418
- volume groups
 - creating, 349–350, 352
 - defined, 572
 - resizing, 353–354
- volume-managing file systems, 356

W

- want, 573
- wants, 573, 390
- wc command, 84, 88–89

- web servers. *See* Apache servers
- while loops, 435–436
- wildcards
 - in automount, 523
 - defined, 573
 - file management, 61–62
 - regular expressions, 91–92
- words, counting, 88–89
- write permissions, 149–152

X

- X forwarding, 111
- XFS
 - defined, 573
 - property management, 329
- xfstool command, 329

Y

- YUM (Yellowdog Update, Modified)
 - defined, 573, 200
 - repositories, role of, 200–201
- yum command
 - common tasks, 209–210
 - dnf command and, 200
 - modules
 - enabling, 223
 - information about, 223
 - installing, 224
 - listing, 221–223
 - managing, 221–225
 - packages
 - finding, 210–211
 - groups of, 216–218
 - history of, 218–219
 - information about, 211–212
 - installing, 212–213
 - listing, 214–215
 - managing, 219
 - removing, 213–214
 - updating, 215–216
- yum group command, 216
- yum groupinstall command, 216

- yum groups info command, 217–218
- yum groups list command, 216–217
- yum history command, 218–219
- yum info command, 211–212
- yum install command, 212–213, 225
- yum install kernel command, 381
- yum list command, 214–215
- yum localinstall command, 225
- yum module enable command, 223
- yum module info command, 223
- yum module install command, 224
- yum module list command, 221–223
- yum provides command, 211

- yum remove command, 213–214
- yum repolist command, 209
- yum search command, 210–211
- yum update command, 215–216
- yum upgrade kernel command, 381
- yum whatprovides command, 211

Z

- zombie processes, 247
- zombies, 573
- zones
 - defined, 573
 - in firewalld, 501

This page intentionally left blank



Video Course Companion to *Red Hat RHCSA 8 Cert Guide*



livelessons®
video instruction from technology experts

Red Hat Certified System Administrator (RHCSA) Complete Video Course, 3rd Edition has more than 16 hours of comprehensive video training—including whiteboard concept teaching, live CLI work, screencast teaching, and hands-on labs—so you have everything you need to prepare for RHCSA success. This new edition of the best-selling *RHCSA Complete Video Course* is fully updated for the RHEL 8 version of the exam. The combination of video, labs, practice exams, and virtual machines is a unique offering that gives you a full toolkit to learn and excel on your exam.

- Module 1: Performing Basic System Management Tasks
- Module 2: Operating Running Systems
- Module 3: Performing Advanced System Administration Tasks
- Module 4: Managing Network Services
- Module 5: Sample Exam

Save 70%*—Use coupon code RHCSA-CVC-70
[pearsonITcertification.com/rhcsa/video](https://t.me/learningnets)

*Discount code RHCSA-CVC-70 confers a 70% discount off the list price of featured video when purchased on [pearsonITcertification.com](https://t.me/learningnets). Offer is subject to change.



PEARSON
IT CERTIFICATION

RHCSA Practice Exam C

General Notes

Here are some tips to ensure your exam starts with a clean environment:

- You do not need any external servers or resources.
 - Do *not* register or connect to any external repositories.
 - Install a new VM according to the instructions in each practice exam.
 - No sample solutions are provided for these practice exams. On the real exam, you need to be able to verify the solutions for yourself as well.
 - You should be able to complete each exam within two hours.
1. Install a RHEL 8 or CentOS 8 virtual machine that meets the following requirements:
 - 2 GB of RAM
 - 20 GB of disk space using default partitioning
 - One additional 20-GB disk that does not have any partitions installed
 - Server with GUI installation pattern
 2. Create user **student** with password **password**, and user **root** with password **password**.
 3. Configure your system to automatically loop-mount the ISO of the installation disk on the directory **/repo**. Configure your system to remove this loop-mounted ISO as the only repository that is used for installation. Do *not* register your system with **subscription-manager**, and remove all reference to external repositories that may already exist.
 4. Reboot your server. Assume that you don't know the root password, and use the appropriate mode to enter a root shell that doesn't require a password. Set the root password to **mypassword**.
 5. Set default values for new users. Make sure that any new user password has a length of at least six characters and must be used for at least three days before it can be reset.

6. Create users **edwin** and **santos** and make them members of the group **sales** as a secondary group membership. Also, create users **serene** and **alex** and make them members of the group **account** as a secondary group.
7. Create shared group directories **/groups/sales** and **/groups/account**, and make sure these groups meet the following requirements:
 - Members of the group **sales** have full access to their directory.
 - Members of the group **account** have full access to their directory.
 - Users have permissions to delete only their own files, but **Alex** is the general manager, so user **alex** has access to delete all users' files.
8. Create a 4-GiB volume group, using a physical extent size of 2 MiB. In this volume group, create a 1-GiB logical volume with the name **myfiles** and mount it persistently on **/myfiles**.
9. Create a group **sysadmins**. Make users **edwin** and **santos** members of this group and ensure that all members of this group can run all administrative commands using **sudo**.
10. Optimize your server with the appropriate profile that optimizes throughput.
11. Add a new disk to your virtual machine with a size of 10 GiB. On this disk, create a VDO volume with a size of 50 GiB and mount it persistently.
12. Configure your server to synchronize time with **serverabc.example.com**, where **serverabc** is an alias to **myserver.example.com**. Note that this server does not have to exist to accomplish this exercise.
13. Configure a web server to use the non-default document root **/webfiles**. In this directory, create a file **index.html** that has the contents **hello world** and then test that it works.

RHCSA Practice Exam D

General Notes

Here are some tips to ensure your exam starts with a clean environment:

- You do not need any external servers or resources.
 - Do *not* register or connect to any external repositories.
 - Install a new VM according to the instructions in each practice exam.
 - No sample solutions are provided for these practice exams. On the real exam, you need to be able to verify the solutions for yourself as well.
 - You should be able to complete each exam within two hours.
1. Install a RHEL 8 or CentOS 8 virtual machine that meets the following requirements:
 - 2 GB of RAM
 - 20 GB of disk space using default partitioning
 - One additional 20-GB disk that does not have any partitions installed
 - Server with GUI installation pattern
 2. Create user **student** with password **password**, and user **root** with password **password**.
 3. Configure your system to automatically loop-mount the ISO of the installation disk on the directory **/repo**. Configure your system to remove this loop-mounted ISO as the only repository that is used for installation. Do *not* register your system with **subscription-manager**, and remove all reference to external repositories that may already exist.
 4. Create a 500-MiB partition on your second hard disk, and format it with the Ext4 file system. Mount it persistently on the directory **/mydata**, using the label **mydata**.
 5. Set default values for new users. A user should get a warning three days before expiration of the current password. Also, new passwords should have a maximum lifetime of 120 days.

6. Create users **edwin** and **santos** and make them members of the group **livingopensource** as a secondary group membership. Also, create users **serene** and **alex** and make them members of the group **operations** as a secondary group.
7. Create shared group directories **/groups/livingopensource** and **/groups/operations**, and make sure the groups meet the following requirements:
 - Members of the group **livingopensource** have full access to their directory.
 - Members of the group **operations** have full access to their directory.
 - Others has no access to any of the directories.
 - Alex is general manager, so user **alex** has read access to all files in both directories and has permissions to delete all files that are created in both directories.
8. Create a 1-GiB swap partition and mount it persistently.
9. Find all files that have the SUID permission set, and write the result to the file **/root/suidfiles**.
10. Create a 1-GiB LVM volume group. In this volume group, create a 512-MiB swap volume and mount it persistently.
11. Add a 10-GiB disk to your virtual machine. On this disk, create a Stratis pool and volume. Use the name **stratisvol** for the volume, and mount it persistently on the directory **/stratis**.
12. Install a web server and configure it to listen on port 8080.
13. Create a configuration that allows user **edwin** to run all administrative commands using **sudo**.

Memory Tables

Chapter 2

Table 2-2 Standard Input, Output, and Error Overview

Name	Default Destination	Use in Redirection	File Descriptor Number
	Computer keyboard	< (same as 0<)	0
	Computer monitor	> (same as 1>)	1
STDERR		2>	2

Table 2-3 Common Bash Redirectors

Redirector	Explanation
(same as 1>)	Redirects STDOUT. If redirection is to a file, the current contents of that file are overwritten.
(same as 1>>)	Redirects STDOUT. If output is written to a file, the output is appended to that file.
	Redirects STDERR.
	Redirects STDERR to the same destination as STDOUT. Notice that this has to be used in combination with normal output redirection, as in ls whuhui > errout 2>&1 .
(same as 0<)	Redirects STDIN.

Table 2-4 vim Essential Commands

vim Command	Explanation
	Switches from input mode to command mode. Press this key before typing any command.
	Switches from command mode to input mode at (i) or after (a) the current cursor position.
	Opens a new line below the current cursor position and goes to input mode.
	Writes the current file and quits.
	Quits the file without applying any changes. The ! forces the command to do its work. Only add the ! if you really know what you are doing.
	Writes the current file with a new filename.
	Deletes the current line.
	Copies the current line.
	Pastes the current selection.
	Enters visual mode, which allows you to select a block of text using the arrow keys. Use d to cut the selection or y to copy it.
	Undoes the last command. Repeat as often as necessary.
	Redoes the last undo.
	Goes to the first line in the document.
	Goes to the last line in the document.
	Searches for text from the current cursor position forward.
	Searches for text from the current cursor position backward.
	Goes to the first position in the current line.
	Goes to the last position in the current line.
	Adds the output of !s (or any other command) in the current file.
	Replaces all occurrences of old with new.

Chapter 3

Table 3-2 FHS Overview

Directory	Use
/	
/boot	

Directory	Use
/dev	
/etc	
/home	
/media, /mnt	
/opt	
/proc	
/root	
/run	
/srv	
/sys	
/tmp	
/usr	
/var	

Chapter 4

Table 4-2 Essential Tools for Managing Text File Contents

Command	Explanation
	Opens the text file in a pager, which allows for easy reading of the text file
	Dumps the contents of the text file on the screen
	Shows the first ten lines of the text file
	Shows the last ten lines of the text file
	Used to filter specific columns or characters from a text file
	Sorts the contents of a text file
	Counts the number of lines, words, and characters in a file

Table 4-3 Most Significant Regular Expressions

Regular Expression	Use
	Matches line that starts with specified text.
	Matches line that ends with specified text.
	Wildcard. (Matches any single character.)
	Matches a, b, or c.
	Matches zero to an infinite number of the previous character.
	Matches exactly two of the previous character.
	Matches a minimum of one and a maximum of three of the previous character.
	Matches zero or one of the previous character. This makes the previous character optional, which in this example would match both <i>color</i> and <i>colour</i> .

Table 4-4 Most Useful **grep** Options

Option	Use
	Not case sensitive. Matches upper- and lowercase letters.
	Only shows lines that do <i>not</i> contain the regular expression.
	Searches files in the current directory and all subdirectories.
	Searches for lines matching more than one regular expression.
	Shows <number> of lines after the matching regular expression.
	Shows <number> of lines before the matching regular expression.

Chapter 5

Table 5.2 Common **ssh** Options

Option	Use
	Verbose; shows in detail what is happening while establishing the connection
	Enables support for graphical applications
	Used to connect to an SSH service that is not listening on the default port 22

Table 5.3 Common **rsync** Options

Option	Use
	Synchronizes the entire directory tree
	Also synchronizes symbolic links
	Preserves symbolic links
	Performs only a dry run, not actually synchronizing anything
	Uses archive mode, thus ensuring that entire subdirectory trees and all file properties will be synchronized
	Uses archive mode, and in addition synchronizes ACLs
	Synchronizes SELinux context as well

Chapter 6

Table 6.2 Methods to Run Tasks with Elevated Permissions

Method	Description
	Opens a subshell as a different user, with the advantage that commands are executed as root only in the subshell
	Enables you to set up an environment where specific tasks are executed with administrative privileges
	Enables you to set up graphical utilities to run with administrative privileges

Chapter 7

Table 7.2 Use of Read, Write, and Execute Permissions

Permission	Applied to Files	Applied to Directories
Read		
Write		
Execute		

Table 7.3 Numeric Representation of Permissions

Permission	Numeric Representation
Read	
Write	
Execute	

Table 7.4 Working with SUID, SGID, and Sticky Bit

Permission	Numeric Value	Relative Value	On Files	On Directories
SUID			User executes file with permissions of file owner.	No meaning.
SGID			User executes file with permissions of group owner.	Files created in directory get the same group owner.
Sticky bit			No meaning.	Prevents users from deleting files from other users.

Table 7.5 umask Values and Their Result

Value	Applied to Files	Applied to Directories
0		Everything
1	Read and write	
2	Read	
3		Read
4	Write	
5	Write	
6	Nothing	
7		Nothing

Chapter 8

Table 8.2 Binary-Decimal Conversion Overview

Binary Value	Decimal Value
	0
	32
	64
	96
	128
	160
	192
	224

Chapter 9

Table 9.2 Key Options in .repo Files

Option	Explanation
	The label used as an identifier in the repository file.
	Specifies the name of the repository.
	Refers to a URL where information about mirror servers for this server can be obtained. Typically used for big online repositories only.
	Refers to the base URL where the RPM packages are found.
	Set to 1 if a GNU Privacy Guard (GPG) integrity check needs to be performed on the packages. If set to 1, a GPG key is required.
	Specifies the location of the GPG key that is used to check package integrity.

Table 9.3 Common **yum** Tasks

Task	Explanation
	Search for the exact name of a package
	Perform a deep search in the package to look for specific files within the package
	Provide more information about the package
	Install the package
	Remove the package
	List all or installed packages
	List package groups
	Install all packages from a group
	Update packages specified
	Remove all stored metadata

Table 9.4 YUM Module Terminology

Item	Explanation
	The default package format. Contains files, as well as metadata that describes how to install the files. Optionally may contain pre- and post-installation scripts as well.
	A delivery mechanism to install RPM packages. In a module different versions and profiles can be provided.
	A specific version of the module.
	A collection of packages that are installed together for a particular use case.

Table 9.5 Common RPM Query Commands

Command	Use
	Uses a filename as its argument to find the specific RPM package a file belongs to.
	Uses the RPM database to provide a list of files in the RPM package.
	Uses the RPM database to provide package information (equivalent to yum info).
	Uses the RPM database to show all documentation that is available in the package.
	Uses the RPM database to show all configuration files that are available in the package.
	Uses the RPM database to show scripts that are used in the package. This is particularly useful if combined with the -p option.
	The -p option is used with all the previously listed options to query individual RPM package files instead of the RPM package database. Using this option before installation helps you find out what is actually in the package before it is installed.
	Shows dependencies for a specific package.
	Shows which parts of a specific package have been changed since installation.
	Verifies all installed packages and shows which parts of the package have been changed since installation. This is an easy and convenient way to do a package integrity check.
	Lists all packages that are installed on this server.

Chapter 10

Table 10.2 Job Management Overview

Command	Use
	Starts the command immediately in the background.
	Stops the job temporarily so that it can be managed. For instance, it can be moved to the background.
	Sends the End Of File (EOF) character to the current job to indicate that it should stop waiting for further input.
	Can be used to cancel the current interactive job.
	Continues the job that has just been frozen using Ctrl-Z in the background.
	Brings back to the foreground the last job that was moved to background execution.
	Shows which jobs are currently running from this shell. Displays job numbers that can be used as an argument to the commands bg and fg .

Table 10.3 Linux Process States Overview

State	Meaning
	The process is currently active and using CPU time, or in the queue of runnable processes waiting to get services.
	The process is waiting for an event to complete.
	The process is in a sleep state that cannot be stopped. This usually happens while a process is waiting for I/O.
	The process has been stopped, which typically has happened to an interactive shell process, using the Ctrl-Z key sequence.
	The process has been stopped but could not be removed by its parent, which has put it in an unmanageable state.

Table 10.4 Tuned Profile Overview

Profile	Use
	The best compromise between power usage and performance
	Based on the balanced profile, but tuned for better response to interactive applications
	Tuned for maximum throughput
	Based on latency-performance, but with additional options to reduce network latency
	Based on throughput-performance, optimizes older CPUs for streaming content
	Tunes for maximum power saving
	Tunes for maximum throughput
	Optimizes Linux for running as a virtual machine
	Optimizes Linux for use as a KVM host

Chapter 11

Table 11.2 Systemd Status Overview

Status	Description
	The unit file has been processed and the unit is active.
	The unit is running with one or more active processes.
	The unit has successfully completed a one-time run.
	The unit is running and waiting for an event.
	The unit is not running.

Status	Description
	The unit will be started at boot time.
	The unit will not be started at boot time.
	The unit cannot be enabled but may be started by another unit automatically.

Table 11.3 `systemctl` Unit Overview Commands

Command	Description
	Shows only service units
	Shows all active service units (same result as the previous command)
	Shows inactive service units as well as active service units
	Shows all services that have failed
	Shows detailed status information about services

Chapter 13

Table 13.2 System Log Files Overview

Log File	Explanation
	The most commonly used log file, it is the generic log file where most messages are written to.
	Contains kernel log messages.
	Contains authentication-related messages. Look here to see which authentication errors have occurred on a server.
	Look here for messages that are related to system startup.
	Contains audit messages. SELinux writes to this file.
	Look here for mail-related messages.
	Provides log files for the Samba service. Notice that Samba by default is not managed through rsyslog, but writes directly to the <code>/var/log</code> directory.
	Contains messages that have been written by the <code>sssd</code> service, which plays an important role in the authentication process.
	Contains log messages that were generated by the print service CUPS.
	Directory that contains log files that are written by the Apache web server. Notice that Apache writes messages to these files directly and not through rsyslog.

Table 13.3 rsyslogd Facilities

Facility	Used by
	Messages related to authentication.
	Messages generated by the crond service.
	Generic facility that can be used for nonspecified daemons.
	Kernel messages.
	Messages generated through the legacy lpd print system.
	Email-related messages.
	Special facility that can be used to write a marker periodically.
	Messages generated by the NNTP news system.
	Same as auth/authpriv. Should not be used anymore.
	Messages generated by the syslog system.
	Messages generated in user space.
	Messages generated by the legacy UUCP system.
	Messages generated by services that are configured by any of the local0 through local7 facilities.

Table 13.4 rsyslogd Priorities

Priority	Description
	Debug messages that will give as much information as possible about service operation.
	Informational messages about normal service operation.
	Informational messages about items that might become an issue later.
	Something is suboptimal, but there is no real error yet.
	A noncritical error has occurred.
	A critical error has occurred.
	Used when the availability of the service is about to be discontinued.
	Message generated when the availability of the service is discontinued.

Chapter 14

Table 14.3 Common Disk Device Types

Device Name	Description
/dev/sda	A hard disk that uses the SCSI driver. Used for SCSI and SATA disk devices. Common on physical servers but also in VMware virtual machines.
	The first hard disk on an NVM Express (NVMe) interface. NVMe is a server-grade method to address advanced SSD devices. Note at the end of the device name that the first disk in this case is referred to as <i>n1</i> instead of <i>a</i> (as in common with the other types).
	The (legacy) IDE disk device type. You will seldom see this device type on modern computers.
	A disk in a KVM virtual machine that uses the virtio disk driver. This is the common disk device type for KVM virtual machines.
	A disk in a Xen virtual machine that uses the Xen virtual disk driver. You see this when installing RHEL as a virtual machine in Xen virtualization. RHEL 8 cannot be used as a Xen hypervisor, but you might see RHEL 8 virtual machines on top of the Xen hypervisor using these disk types.

Table 14.4 File System Overview

File System	Description
	The default file system in RHEL 8.
	The default file system in previous versions of RHEL; still available and supported in RHEL 8.
	The previous version of Ext4. On RHEL 8, there is no need to use Ext3 anymore.
	A very basic file system that was developed in the early 1990s. There is no need to use this file system on RHEL 8 anymore.
	A relatively new file system that is not supported in RHEL 8.
	A Windows-compatible file system that is not supported on RHEL 8.
	A file system that offers compatibility with Windows and Mac and is the functional equivalent of the FAT32 file system. Useful to use on USB thumb drives that are used to exchange data with other computers but not on a server's hard disks.

Table 14.5 /etc/fstab Fields

Field	Description
	The device that must be mounted. A device name, UUID, or label can be used.
	The directory or kernel interface where the device needs to be mounted.
	The file system type.
	Mount options.
	Use 1 to enable support to back up using the dump utility. This may be necessary for some backup solutions.
	Specifies whether the file system should be checked automatically when booting. Use 0 to disable automated check, 1 if this is the root file system and it has to be checked automatically, and 2 for all other file systems that need automatic checking while booting. Network file systems should have this option set to 0.

Table 14.6 Common Mount Options

Option	Use
	The file system will [not] be mounted automatically.
	Adds support for file system access control lists (see Chapter 7, “Permissions Management”).
	Adds support for user-extended attributes (see Chapter 7).
	Mounts the file system in read-only mode.
	Disables or enables access time modifications.
	Denies or allows execution of program files from the file system.
	Use this to mount a network file system. This tells fstab to wait until the network is available before mounting this file system.

Chapter 15

Table 15.2 LVM Management Essential Commands

Command	Explanation
	Creates physical volumes
	Shows a summary of available physical volumes
	Shows a list of physical volumes and their properties

Command	Explanation
	Removes the physical volume signature from a block device
	Creates volume groups
	Shows a summary of available volume groups
	Shows a detailed list of volume groups and their properties
	Removes a volume group
	Creates logical volumes
	Shows a summary of all available logical volumes
	Shows a detailed list of available logical volumes and their properties
	Removes a logical volume

Chapter 16

Table 16.2 Linux Kernel Module Management Overview

Command	Use
	Lists currently loaded kernel modules
	Displays information about kernel modules
	Loads kernel modules, including all of their dependencies
	Unloads kernel modules, considering kernel module dependencies

Chapter 18

Table 18.2 Boot Phase Configuration and Troubleshooting Overview

Boot Phase	Configuring It	Fixing It
	Hardware configuration (F2, Esc, F10, or another key).	Replace hardware.
	BIOS/UEFI configuration or hardware boot menu.	Replace hardware or use rescue system.
	grub2-install and edits to <code>/etc/defaults/grub</code> .	GRUB boot prompt and edits to <code>/etc/defaults/grub</code> , followed by grub2-mkconfig .
	Edits to the GRUB configuration and <code>/etc/dracut.conf</code> .	GRUB boot prompt and edits to <code>/etc/defaults/grub</code> , followed by <code>grub2-mkconfig</code> .

Boot Phase	Configuring It	Fixing It
	Compiled into initramfs.	init= kernel boot argument, rd.break kernel boot argument.
	Compiled into initramfs.	Not typically required.
	Edits to the /etc/fstab file.	Edits to the /etc/fstab file.
	Using systemctl set-default to create the /etc/systemd/system/default.target symbolic link	Start the rescue.target as a kernel boot argument.

Chapter 20

Table 20.2 Most Useful sshd Configuration Options

Option	Use
	Defines the TCP listening port.
	Indicates whether to allow or disallow root login.
	Specifies the maximum number of authentication tries. After reaching half of this number, failures are logged to syslog.
	Indicates the maximum number of sessions that can be open from one IP address.
	Specifies a space-separated list of users who are allowed to connect to the server.
	Specifies whether to allow password authentication. This option is on by default.
	Indicates whether authentication through the GSSAPI needs to be enabled. Used for Kerberos-based authentication.
	Specifies whether or not to clean up inactive TCP connections.
	Specifies the interval, in seconds, that packets are sent to the client to figure out if the client is still alive.
	Specifies the number of client alive packets that need to be sent.
	If on, uses DNS name lookup to match incoming IP addresses to names.
	Specifies the interval, in seconds, that a client sends a packet to a server to keep connections alive.
	Specifies the maximum number of packets a client sends to a server to keep connections alive.

Chapter 22

Table 22.2 SELinux Core Elements

Element	Use
	A collection of rules that define which source has access to which target.
	The object that is trying to access a target. Typically a user or a process.
	The thing that a source domain is trying to access. Typically a file or a port.
	A security label that is used to categorize objects in SELinux.
	A specific part of the policy that determines which source domain has which access permissions to which target domain.
	Same as a context label, defined to determine which source domain has access to which target domain.

Chapter 23

Table 23.2 FirewallD Default Zones

Zone Name	Default Settings
	Incoming network connections are rejected with an “icmp-host-prohibited” message. Only network connections that were initiated on this system are allowed.
	For use on computers in the demilitarized zone. Only selected incoming connections are accepted, and limited access to the internal network is allowed.
	Any incoming packets are dropped and there is no reply.
	For use on external networks with masquerading (Network Address Translation [NAT]) enabled, used especially on routers. Only selected incoming connections are accepted.
	For use with home networks. Most computers on the same network are trusted, and only selected incoming connections are accepted.
	For use in internal networks. Most computers on the same network are trusted, and only selected incoming connections are accepted.
	For use in public areas. Other computers in the same network are not trusted, and limited connections are accepted. This is the default zone for all newly created network interfaces.
	All network connections are accepted.
	For use in work areas. Most computers on the same network are trusted, and only selected incoming connections are accepted.

Table 23.3 Common `firewall-cmd` Options

firewall-cmd Options	Explanation
	Lists all available zones
	Shows the zone currently set as default zone
	Changes the default zone
	Shows all available services
	Shows services currently in use
	Adds a service to the current default zone or the zone that is specified
	Removes a service from the configuration
	Lists all configurations in a zone
	Adds a port and protocol
	Removes a port from the configuration
	Adds an interface to the default zone or a specific zone that is specified
	Removes an interface from a specific zone
	Adds a specific IP address
	Removes an IP address from the configuration
	Writes configuration to disk and not to runtime
	Reloads the on-disk configuration

Chapter 25

Table 25.2 Understanding Linux Time

Concept	Explanation
	The hardware clock that resides on the main card of a computer system
	Same as the hardware clock
	The time that is maintained by the operating system
	Similar to system time

Concept	Explanation
	A worldwide standard time
	Calculation that is made to change time automatically when DST changes occur
	The time that corresponds to the time in the current time zone

Table 25.3 Commands Related to RHEL 8 Time Management

Command	Short Description
	Manages local time
	Manages hardware time
	Developed to manage all aspects of time on RHEL 8

Table 25.4 `timedatectl` Command Overview

Command	Explanation
	Shows current time settings
	Sets the current time
	Sets the current time zone
	Shows a list of all time zones
	Controls whether the RTC (the real-time clock, normally referred to as the hardware clock) is in local time
	Controls whether NTP is enabled

Memory Tables Answer Key

Chapter 2

Table 2-2 Standard Input, Output, and Error Overview

Name	Default Destination	Use in Redirection	File Descriptor Number
STDIN	Computer keyboard	< (same as 0<)	0
STDOUT	Computer monitor	> (same as 1>)	1
STDERR	Computer monitor	2>	2

Table 2-3 Common Bash Redirectors

Redirector	Explanation
> (same as 1>)	Redirects STDOUT. If redirection is to a file, the current contents of that file are overwritten.
>> (same as 1>>)	Redirects STDOUT. If output is written to a file, the output is appended to that file.
2>	Redirects STDERR.
2>&1	Redirects STDERR to the same destination as STDOUT. Notice that this has to be used in combination with normal output redirection, as in ls whuhiu > errout 2>&1 .
< (same as 0<)	Redirects STDIN.

Table 2-4 vim Essential Commands

vim Command	Explanation
Esc	Switches from input mode to command mode. Press this key before typing any command.
i, a	Switches from command mode to input mode at (i) or after (a) the current cursor position.
o	Opens a new line below the current cursor position and goes to input mode.
:wq	Writes the current file and quits.
:q!	Quits the file without applying any changes. The ! forces the command to do its work. Only add the ! if you really know what you are doing.
:w filename	Writes the current file with a new filename.
dd	Deletes the current line.
yy	Copies the current line.
p	Pastes the current selection.
v	Enters visual mode, which allows you to select a block of text using the arrow keys. Use d to cut the selection or y to copy it.
u	Undoes the last command. Repeat as often as necessary.
Ctrl-r	Redoes the last undo.
gg	Goes to the first line in the document.
G	Goes to the last line in the document.
/text	Searches for <i>text</i> from the current cursor position forward.
?text	Searches for <i>text</i> from the current cursor position backward.
^	Goes to the first position in the current line.
\$	Goes to the last position in the current line.
!ls	Adds the output of ls (or any other command) in the current file.
:%s/old/new/g	Replaces all occurrences of <i>old</i> with <i>new</i> .

Chapter 3

Table 3-2 FHS Overview

Directory	Use
/	The root directory. This is where the file system tree starts.
/boot	Contains all files and directories that are needed to boot the Linux kernel.

Directory	Use
/dev	Contains device files that are used for accessing physical devices. This directory is essential during boot.
/etc	Contains configuration files that are used by programs and services on your server. This directory is essential during boot.
/home	Used for local user home directories.
/media, /mnt	Contain directories that are used for mounting devices in the file system tree.
/opt	Used for optional packages that may be installed on your server.
/proc	Used by the proc file system. This is a file system structure that gives access to kernel information.
/root	The home directory of the root user.
/run	Contains process and user-specific information that has been created since the last boot.
/srv	May be used for data by services like NFS, FTP, and HTTP.
/sys	Used as an interface to different hardware devices that is managed by the Linux kernel and associated processes.
/tmp	Contains temporary files that may be deleted without any warning during boot.
/usr	Contains subdirectories with program files, libraries for these program files, and documentation about them.
/var	Contains files that may change in size dynamically, such as log files, mail boxes, and spool files.

Chapter 4

Table 4-2 Essential Tools for Managing Text File Contents

Command	Explanation
less	Opens the text file in a pager, which allows for easy reading
cat	Dumps the contents of the text file on the screen
head	Shows the first ten lines of the text file
tail	Shows the last ten lines of the text file
cut	Used to filter specific columns or characters from a text file
sort	Sorts contents of a text file
wc	Counts the number of lines, words, and characters in a file

Table 4-3 Most Significant Regular Expressions

Regular Expression	Use
<code>^text</code>	Matches line that starts with specified text.
<code>text\$</code>	Matches line that ends with specified text.
<code>.</code>	Wildcard. (Matches any single character.)
<code>[abc]</code>	Matches a, b, or c.
<code>*</code>	Matches zero to an infinite number of the previous character.
<code>\{2\}</code>	Matches exactly two of the previous character.
<code>\{1,3\}</code>	Matches a minimum of one and a maximum of three of the previous character.
<code>colou?r</code>	Matches zero or one of the previous character. This makes the previous character optional, which in this example would match both <i>color</i> and <i>colour</i> .

Table 4-4 Most Useful **grep** Options

Option	Use
<code>-i</code>	Not case sensitive. Matches upper- and lowercase letters.
<code>-v</code>	Only shows lines that do <i>not</i> contain the regular expression.
<code>-r</code>	Searches files in the current directory and all subdirectories.
<code>-e</code>	Searches for lines matching more than one regular expression.
<code>-A <number></code>	Shows <number> of lines after the matching regular expression.
<code>-B <number></code>	Shows <number> of lines before the matching regular expression.

Chapter 5

Table 5.2 Common **ssh** Options

Option	Use
<code>-v</code>	Verbose; shows in detail what is happening while establishing the connection
<code>-Y</code>	Enables support for graphical applications
<code>-p <PORT></code>	Used to connect to an SSH service that is not listening on the default port 22

Table 5.3 Common **rsync** Options

Option	Use
-r	Synchronizes the entire directory tree
-l	Also synchronizes symbolic links
-p	Preserves symbolic links
-n	Performs only a dry run, not actually synchronizing anything
-a	Uses archive mode, thus ensuring that entire subdirectory trees and all file properties will be synchronized
-A	Uses archive mode, and in addition synchronizes ACLs
-X	Synchronizes SELinux context as well

Chapter 6

Table 6.2 Methods to Run Tasks with Elevated Permissions

Method	Description
su	Opens a subshell as a different user, with the advantage that commands are executed as root only in the subshell
sudo	Enables you to set up an environment where specific tasks are executed with administrative privileges
PolicyKit	Enables you to set up graphical utilities to run with administrative privileges

Chapter 7

Table 7.2 Use of Read, Write, and Execute Permissions

Permission	Applied to Files	Applied to Directories
Read	Open a file	List contents of directory
Write	Change contents of a file	Create and delete files
Execute	Run a program file	Change to the directory

Table 7.3 Numeric Representation of Permissions

Permission	Numeric Representation
Read	4
Write	2
Execute	1

Table 7.4 Working with SUID, SGID, and Sticky Bit

Permission	Numeric Value	Relative Value	On Files	On Directories
SUID	4	u+s	User executes file with permissions of file owner.	No meaning.
SGID	2	g+s	User executes file with permissions of group owner.	Files created in directory get the same group owner.
Sticky bit	1	+t	No meaning.	Prevents users from deleting files from other users.

Table 7.5 umask Values and Their Result

Value	Applied to Files	Applied to Directories
0	Read and write	Everything
1	Read and write	Read and write
2	Read	Read and execute
3	Read	Read
4	Write	Write and execute
5	Write	Write
6	Nothing	Execute
7	Nothing	Nothing

Chapter 8

Table 8.2 Binary-Decimal Conversion Overview

Binary Value	Decimal Value
00000000	0
00100000	32
01000000	64
01100000	96
10000000	128
10100000	160
11000000	192
11100000	224

Chapter 9

Table 9.2 Key Options in .repo Files

Option	Explanation
[label]	The label used as an identifier in the repository file.
name=	Specifies the name of the repository.
mirrorlist=	Refers to a URL where information about mirror servers for this server can be obtained. Typically used for big online repositories only.
baseurl=	Refers to the base URL where the RPM packages are found.
gpgcheck=	Set to 1 if a GNU Privacy Guard (GPG) integrity check needs to be performed on the packages. If set to 1, a GPG key is required.
gpgkey=	Specifies the location of the GPG key that is used to check package integrity.

Table 9.3 Common **yum** Tasks

Task	Explanation
search	Search for the exact name of a package
[what]provides */name	Perform a deep search in the package to look for specific files within the package
info	Provide more information about the package
install	Install the package
remove	Remove the package
list [all installed]	List all or installed packages
group list	List package groups
group install	Install all packages from a group
update	Update packages specified
clean all	Remove all stored metadata

Table 9.4 YUM Module Terminology

Item	Explanation
RPM	The default package format. Contains files, as well as metadata that describes how to install the files. Optionally may contain pre- and post-installation scripts as well.
Module	A delivery mechanism to install RPM packages. In a module different versions and profiles can be provided.
Application stream	A specific version of the module.
Profile	A collection of packages that are installed together for a particular use case.

Table 9.5 Common RPM Query Commands

Command	Use
rpm -qf	Uses a filename as its argument to find the specific RPM package a file belongs to.
rpm -ql	Uses the RPM database to provide a list of files in the RPM package.
rpm -qi	Uses the RPM database to provide package information (equivalent to yum info).
rpm -qd	Uses the RPM database to show all documentation that is available in the package.
rpm -qc	Uses the RPM database to show all configuration files that are available in the package.
rpm -q --scripts	Uses the RPM database to show scripts that are used in the package. This is particularly useful if combined with the -p option.
rpm -qp <pkg>	The -p option is used with all the previously listed options to query individual RPM package files instead of the RPM package database. Using this option before installation helps you find out what is actually in the package before it is installed.
rpm -qR	Shows dependencies for a specific package.
rpm -V	Shows which parts of a specific package have been changed since installation.
rpm -Va	Verifies all installed packages and shows which parts of the package have been changed since installation. This is an easy and convenient way to do a package integrity check.
rpm -qa	Lists all packages that are installed on this server.

Chapter 10

Table 10.2 Job Management Overview

Command	Use
& (used at the end of a command line)	Starts the command immediately in the background.
Ctrl-Z	Stops the job temporarily so that it can be managed. For instance, it can be moved to the background.
Ctrl-D	Sends the End Of File (EOF) character to the current job to indicate that it should stop waiting for further input.
Ctrl-C	Can be used to cancel the current interactive job.
bg	Continues the job that has just been frozen using Ctrl-Z in the background.
fg	Brings back to the foreground the last job that was moved to background execution.
jobs	Shows which jobs are currently running from this shell. Displays job numbers that can be used as an argument to the commands bg and fg .

Table 10.3 Linux Process States Overview

State	Meaning
Running (R)	The process is currently active and using CPU time, or in the queue of runnable processes waiting to get services.
Sleeping (S)	The process is waiting for an event to complete.
Uninterruptable sleep (D)	The process is in a sleep state that cannot be stopped. This usually happens while a process is waiting for I/O.
Stopped (T)	The process has been stopped, which typically has happened to an interactive shell process, using the Ctrl-Z key sequence.
Zombie (Z)	The process has been stopped but could not be removed by its parent, which has put it in an unmanageable state.

Table 10.4 Tuned Profile Overview

Profile	Use
balanced	The best compromise between power usage and performance
desktop	Based on the balanced profile, but tuned for better response to interactive applications
latency-performance	Tuned for maximum throughput
network-latency	Based on latency-performance, but with additional options to reduce network latency
network-throughput	Based on throughput-performance, optimizes older CPUs for streaming content
powersave	Tunes for maximum power saving
throughput-performance	Tunes for maximum throughput
virtual-guest	Optimizes Linux for running as a virtual machine
virtual-host	Optimizes Linux for use as a KVM host

Chapter 11

Table 11.2 Systemd Status Overview

Status	Description
Loaded	The unit file has been processed and the unit is active.
Active(running)	The unit is running with one or more active processes.
Active(exited)	The unit has successfully completed a one-time run.
Active(waiting)	The unit is running and waiting for an event.
Inactive(dead)	The unit is not running.

Status	Description
Enabled	The unit will be started at boot time.
Disabled	The unit will not be started at boot time.
Static	The unit cannot be enabled but may be started by another unit automatically.

Table 11.3 `systemctl` Unit Overview Commands

Command	Description
<code>systemctl --type=service</code>	Shows only service units
<code>systemctl list-units --type=service</code>	Shows all active service units (same result as the previous command)
<code>systemctl list-units --type=service --all</code>	Shows inactive service units as well as active service units
<code>systemctl --failed --type=service</code>	Shows all services that have failed
<code>systemctl status -l your.service</code>	Shows detailed status information about services

Chapter 13

Table 13.2 System Log Files Overview

Log File	Explanation
<code>/var/log/messages</code>	The most commonly used log file, it is the generic log file where most messages are written to.
<code>/var/log/dmesg</code>	Contains kernel log messages.
<code>/var/log/secure</code>	Contains authentication-related messages. Look here to see which authentication errors have occurred on a server.
<code>/var/log/boot.log</code>	Look here for messages that are related to system startup.
<code>/var/log/audit/audit.log</code>	Contains audit messages. SELinux writes to this file.
<code>/var/log/maillog</code>	Look here for mail-related messages.
<code>/var/log/samba</code>	Provides log files for the Samba service. Notice that Samba by default is not managed through rsyslog, but writes directly to the <code>/var/log</code> directory.
<code>/var/log/sss</code>	Contains messages that have been written by the sssd service, which plays an important role in the authentication process.
<code>/var/log/cups</code>	Contains log messages that were generated by the print service CUPS.
<code>/var/log/httpd/</code>	Directory that contains log files that are written by the Apache web server. Notice that Apache writes messages to these files directly and not through rsyslog.

Table 13.3 rsyslogd Facilities

Facility	Used by
auth / authpriv	Messages related to authentication.
cron	Messages generated by the crond service.
daemon	Generic facility that can be used for nonspecified daemons.
kern	Kernel messages.
lpr	Messages generated through the legacy lpd print system.
mail	Email-related messages.
mark	Special facility that can be used to write a marker periodically.
news	Messages generated by the NNTP news system.
security	Same as auth/authpriv. Should not be used anymore.
syslog	Messages generated by the syslog system.
user	Messages generated in user space.
uucp	Messages generated by the legacy UUCP system.
local0-7	Messages generated by services that are configured by any of the local0 through local7 facilities.

Table 13.4 rsyslogd Priorities

Priority	Description
debug	Debug messages that will give as much information as possible about service operation.
info	Informational messages about normal service operation.
notice	Informational messages about items that might become an issue later.
warning / warn	Something is suboptimal, but there is no real error yet.
err / error	A noncritical error has occurred.
crit	A critical error has occurred.
alert	Used when the availability of the service is about to be discontinued.
emerg / panic	Message generated when the availability of the service is discontinued.

Chapter 14

Table 14.3 Common Disk Device Types

Device Name	Description
/dev/sda	A hard disk that uses the SCSI driver. Used for SCSI and SATA disk devices. Common on physical servers but also in VMware virtual machines.
/dev/nvme0n1	The first hard disk on an NVM Express (NVMe) interface. NVMe is a server-grade method to address advanced SSD devices. Note at the end of the device name that the first disk in this case is referred to as <i>n1</i> instead of <i>a</i> (as in common with the other types).
/dev/hda	The (legacy) IDE disk device type. You will seldom see this device type on modern computers.
/dev/sda	A disk in a KVM virtual machine that uses the virtio disk driver. This is the common disk device type for KVM virtual machines.
/dev/xvda	A disk in a Xen virtual machine that uses the Xen virtual disk driver. You see this when installing RHEL as a virtual machine in Xen virtualization. RHEL 8 cannot be used as a Xen hypervisor, but you might see RHEL 8 virtual machines on top of the Xen hypervisor using these disk types.

Table 14.4 File System Overview

File System	Description
XFS	The default file system in RHEL 8.
Ext4	The default file system in previous versions of RHEL; still available and supported in RHEL 8.
Ext3	The previous version of Ext4. On RHEL 8, there is no need to use Ext3 anymore.
Ext2	A very basic file system that was developed in the early 1990s. There is no need to use this file system on RHEL 8 anymore.
Btrfs	A relatively new file system that is not supported in RHEL 8.
NTFS	A Windows-compatible file system that is not supported on RHEL 8.
VFAT	A file system that offers compatibility with Windows and Mac and is the functional equivalent of the FAT32 file system. Useful on USB thumb drives that are used to exchange data with other computers but not on a server's hard disks.

Table 14.5 /etc/fstab Fields

Field	Description
Device	The device that must be mounted. A device name, UUID, or label can be used.
Mount Point	The directory or kernel interface where the device needs to be mounted.
File System	The file system type.
Mount Options	Mount options.
Dump Support	Use 1 to enable support to back up using the dump utility. This may be necessary for some backup solutions.
Automatic Check	Specifies whether the file system should be checked automatically when booting. Use 0 to disable automated check, 1 if this is the root file system and it has to be checked automatically, and 2 for all other file systems that need automatic checking while booting. Network file systems should have this option set to 0.

Table 14.6 Common Mount Options

Option	Use
auto/ noauto	The file system will [not] be mounted automatically.
acl	Adds support for file system access control lists (see Chapter 7, “Permissions Management”).
user_xattr	Adds support for user-extended attributes (see Chapter 7).
ro	Mounts the file system in read-only mode.
atime / noatime	Disables or enables access time modifications.
noexec / exec	Denies or allows execution of program files from the file system.
_netdev	Use this to mount a network file system. This tells fstab to wait until the network is available before mounting this file system.

Chapter 15

Table 15.2 LVM Management Essential Commands

Command	Explanation
pvccreate	Creates physical volumes
pvs	Shows a summary of available physical volumes
pvdisplay	Shows a list of physical volumes and their properties

Command	Explanation
pvremove	Removes the physical volume signature from a block device
vgcreate	Creates volume groups
vgs	Shows a summary of available volume groups
vgdisplay	Shows a detailed list of volume groups and their properties
vgremove	Removes a volume group
lvcreate	Creates logical volumes
lvs	Shows a summary of all available logical volumes
lvdisplay	Shows a detailed list of available logical volumes and their properties
lvremove	Removes a logical volume

Chapter 16

Table 16.2 Linux Kernel Module Management Overview

Command	Use
lsmod	Lists currently loaded kernel modules
modinfo	Displays information about kernel modules
modprobe	Loads kernel modules, including all of their dependencies
modprobe -r	Unloads kernel modules, considering kernel module dependencies

Chapter 18

Table 18.2 Boot Phase Configuration and Troubleshooting Overview

Boot Phase	Configuring It	Fixing It
POST	Hardware configuration (F2, Esc, F10, or another key).	Replace hardware.
Selecting the bootable device	BIOS/UEFI configuration or hardware boot menu.	Replace hardware or use rescue system.
Loading the boot loader	grub2-install and edits to <code>/etc/defaults/grub</code> .	GRUB boot prompt and edits to <code>/etc/defaults/grub</code> , followed by grub2-mkconfig .
Loading the kernel	Edits to the GRUB configuration and <code>/etc/dracut.conf</code> .	GRUB boot prompt and edits to <code>/etc/defaults/grub</code> , followed by grub2-mkconfig .

Boot Phase	Configuring It	Fixing It
Starting /sbin/init	Compiled into initramfs.	init= kernel boot argument, rd.break kernel boot argument.
Processing initrd. target	Compiled into initramfs.	Not typically required.
Switch to the root file system	Edits to the /etc/fstab file.	Edits to the /etc/fstab file.
Running the default target	Using systemctl set-default to create the /etc/systemd/system/default.target symbolic link	Start the rescue.target as a kernel boot argument.

Chapter 20

Table 20.2 Most Useful sshd Configuration Options

Option	Use
Port	Defines the TCP listening port.
PermitRootLogin	Indicates whether to allow or disallow root login.
MaxAuthTries	Specifies the maximum number of authentication tries. After reaching half of this number, failures are logged to syslog.
MaxSessions	Indicates the maximum number of sessions that can be open from one IP address.
AllowUsers	Specifies a space-separated list of users that are allowed to connect to the server.
PasswordAuthentication	Specifies whether to allow password authentication. This option is on by default.
GSSAPIAuthentication	Indicates whether authentication through the GSSAPI needs to be enabled. Used for Kerberos-based authentication.
TCPKeepAlive	Specifies whether or not to clean up inactive TCP connections.
ClientAliveInterval	Specifies the interval, in seconds, that packets are sent to the client to figure out if the client is still alive.
ClientAliveCountMax	Specifies the number of client alive packets that need to be sent.
UseDNS	If on, uses DNS name lookup to match incoming IP addresses to names.
ServerAliveInterval	Specifies the interval, in seconds, that a client sends a packet to a server to keep connections alive.
ServerAliveCountMax	Specifies the maximum number of packets a client sends to a server to keep connections alive.

Chapter 22

Table 22.2 SELinux Core Elements

Element	Use
Policy	A collection of rules that define which source has access to which target.
Source domain	The object that is trying to access a target. Typically a user or a process.
Target domain	The thing that a source domain is trying to access. Typically a file or a port.
Context	A security label that is used to categorize objects in SELinux.
Rule	A specific part of the policy that determines which source domain has which access permissions to which target domain.
Labels	Same as a context label, defined to determine which source domain has access to which target domain.

Chapter 23

Table 23.2 FirewallD Default Zones

Zone Name	Default Settings
block	Incoming network connections are rejected with an “icmp-host-prohibited” message. Only network connections that were initiated on this system are allowed.
dmz	For use on computers in the demilitarized zone. Only selected incoming connections are accepted, and limited access to the internal network is allowed.
drop	Any incoming packets are dropped and there is no reply.
external	For use on external networks with masquerading (Network Address Translation [NAT]) enabled, used especially on routers. Only selected incoming connections are accepted.
home	For use with home networks. Most computers on the same network are trusted, and only selected incoming connections are accepted.
internal	For use in internal networks. Most computers on the same network are trusted, and only selected incoming connections are accepted.
public	For use in public areas. Other computers in the same network are not trusted, and limited connections are accepted. This is the default zone for all newly created network interfaces.
trusted	All network connections are accepted.
work	For use in work areas. Most computers on the same network are trusted, and only selected incoming connections are accepted.

Table 23.3 Common `firewall-cmd` Options

firewall-cmd Options	Explanation
<code>--get-zones</code>	Lists all available zones
<code>--get-default-zone</code>	Shows the zone currently set as the default zone
<code>--set-default-zone=<ZONE></code>	Changes the default zone
<code>--get-services</code>	Shows all available services
<code>--list-services</code>	Shows services currently in use
<code>--add-service=<service-name></code> <code>[--zone=<ZONE>]</code>	Adds a service to the current default zone or the zone that is specified
<code>--remove-service=<service-name></code>	Removes a service from the configuration
<code>--list-all [--zone=<ZONE>]</code>	Lists all configurations in a zone
<code>--add-port=<port/protocol></code> <code>[--zone=<ZONE>]</code>	Adds a port and protocol
<code>--remove-port=<port/protocol></code> <code>[--zone=<ZONE>]</code>	Removes a port from the configuration
<code>--add-interface=<INTERFACE></code> <code>[--zone=<ZONE>]</code>	Adds an interface to the default zone or a specific zone that is specified
<code>--remove-interface=<INTERFACE></code> <code>[--zone=<ZONE>]</code>	Removes an interface from a specific zone
<code>--add-source=<ipaddress/netmask></code> <code>[--zone=<ZONE>]</code>	Adds a specific IP address
<code>--remove-source=<ipaddress/netmask></code> <code>[--zone=<ZONE>]</code>	Removes an IP address from the configuration
<code>--permanent</code>	Writes configuration to disk and not to runtime
<code>--reload</code>	Reloads the on-disk configuration

Chapter 25

Table 25.2 Understanding Linux Time

Concept	Explanation
Hardware clock	The hardware clock that resides on the main card of a computer system
Real-time clock	Same as the hardware clock
System time	The time that is maintained by the operating system
Software clock	Similar to system time

Concept	Explanation
Coordinated Universal Time (UTC)	A worldwide standard time
Daylight saving time	Calculation that is made to change time automatically when DST changes occur
Local time	The time that corresponds to the time in the current time zone

Table 25.3 Commands Related to RHEL 8 Time Management

Command	Short Description
<code>date</code>	Manages local time
<code>hwclock</code>	Manages hardware time
<code>timedatectl</code>	Developed to manage all aspects of time on RHEL 8

Table 25.4 `timedatectl` Command Overview

Command	Explanation
<code>status</code>	Shows the current time settings
<code>set-time TIME</code>	Sets the current time
<code>set-timezone ZONE</code>	Sets the current time zone
<code>list-timezone</code>	Shows a list of all time zones
<code>set-local-rtc [0 1]</code>	Controls whether the RTC (the real-time clock, normally referred to as the hardware clock) is in local time
<code>set-ntp [0 1]</code>	Control whether NTP is enabled

Appendix D

Study Planner

Practice Test	Reading	Task

Element	Task	Goal Date	First Date Completed	Second Date Completed (Optional)	Notes
Introduction	Read Introduction				
1. Installing Red Hat Enterprise Linux Server	Read Foundation Topics				
1. Installing Red Hat Enterprise Linux Server	Review Key Topics				
1. Installing Red Hat Enterprise Linux Server	Define Key Terms				
1. Installing Red Hat Enterprise Linux Server	Answer Chapter Review Questions				
1. Installing Red Hat Enterprise Linux Server	Complete Chapter Lab				
2. Using Essential Tools	Read Foundation Topics				
2. Using Essential Tools	Complete Chapter Exercises				
2. Using Essential Tools	Review Key Topics				
2. Using Essential Tools	Complete Memory Tables				
2. Using Essential Tools	Define Key Terms				
2. Using Essential Tools	Answer Chapter Review Questions				
2. Using Essential Tools	Complete Chapter Lab				
3. Essential File Management Tools	Read Foundation Topics				
3. Essential File Management Tools	Complete Chapter Exercises				
3. Essential File Management Tools	Review Key Topics				
3. Essential File Management Tools	Complete Memory Tables				
3. Essential File Management Tools	Define Key Terms				
3. Essential File Management Tools	Answer Chapter Review Questions				
3. Essential File Management Tools	Complete Chapter Lab				
4. Working with Text Files	Read Foundation Topics				
4. Working with Text Files	Complete Chapter Exercises				

4. Working with Text Files	Review Key Topics				
4. Working with Text Files	Complete Memory Tables				
4. Working with Text Files	Define Key Terms				
4. Working with Text Files	Answer Chapter Review Questions				
4. Working with Text Files	Complete Chapter Lab				
5. Connecting to Red Hat Enterprise Linux 8	Read Foundation Topics				
5. Connecting to Red Hat Enterprise Linux 8	Complete Chapter Exercises				
5. Connecting to Red Hat Enterprise Linux 8	Review Key Topics				
5. Connecting to Red Hat Enterprise Linux 8	Complete Memory Tables				
5. Connecting to Red Hat Enterprise Linux 8	Define Key Terms				
5. Connecting to Red Hat Enterprise Linux 8	Answer Chapter Review Questions				
5. Connecting to Red Hat Enterprise Linux 8	Complete Chapter Lab				
6. User and Group Management	Read Foundation Topics				
6. User and Group Management	Complete Chapter Exercises				
6. User and Group Management	Review Key Topics				
6. User and Group Management	Complete Memory Tables				
6. User and Group Management	Define Key Terms				
6. User and Group Management	Answer Chapter Review Questions				
6. User and Group Management	Complete Chapter Lab				
7. Permissions Management	Read Foundation Topics				
7. Permissions Management	Complete Chapter Exercises				
7. Permissions Management	Review Key Topics				
7. Permissions Management	Complete Memory Tables				
7. Permissions Management	Define Key Terms				
7. Permissions Management	Answer Chapter Review Questions				
7. Permissions Management	Complete Chapter Lab				
8. Configuring Networking	Read Foundation Topics				
8. Configuring Networking	Complete Chapter Exercises				
8. Configuring Networking	Review Key Topics				
8. Configuring Networking	Complete Memory Tables				
8. Configuring Networking	Define Key Terms				
8. Configuring Networking	Answer Chapter Review Questions				
8. Configuring Networking	Complete Chapter Lab				

9. Managing Software	Read Foundation Topics				
9. Managing Software	Complete Chapter Exercises				
9. Managing Software	Review Key Topics				
9. Managing Software	Complete Memory Tables				
9. Managing Software	Define Key Terms				
9. Managing Software	Answer Chapter Review Questions				
9. Managing Software	Complete Chapter Lab				
10. Managing Processes	Read Foundation Topics				
10. Managing Processes	Complete Chapter Exercises				
10. Managing Processes	Review Key Topics				
10. Managing Processes	Complete Memory Tables				
10. Managing Processes	Define Key Terms				
10. Managing Processes	Answer Chapter Review Questions				
10. Managing Processes	Complete Chapter Lab				
11. Working with Systemd	Read Foundation Topics				
11. Working with Systemd	Complete Chapter Exercises				
11. Working with Systemd	Review Key Topics				
11. Working with Systemd	Complete Memory Tables				
11. Working with Systemd	Define Key Terms				
11. Working with Systemd	Answer Chapter Review Questions				
11. Working with Systemd	Complete Chapter Lab				
12. Scheduling Tasks	Read Foundation Topics				
12. Scheduling Tasks	Complete Chapter Exercises				
12. Scheduling Tasks	Review Key Topics				
12. Scheduling Tasks	Define Key Terms				
12. Scheduling Tasks	Answer Chapter Review Questions				
12. Scheduling Tasks	Complete Chapter Lab				
13. Configuring Logging	Read Foundation Topics				
13. Configuring Logging	Complete Chapter Exercises				
13. Configuring Logging	Review Key Topics				
13. Configuring Logging	Complete Memory Tables				
13. Configuring Logging	Define Key Terms				
13. Configuring Logging	Answer Chapter Review Questions				
13. Configuring Logging	Complete Chapter Lab				
14. Managing Storage	Read Foundation Topics				

14. Managing Storage	Complete Chapter Exercises				
14. Managing Storage	Review Key Topics				
14. Managing Storage	Complete Memory Tables				
14. Managing Storage	Define Key Terms				
14. Managing Storage	Answer Chapter Review Questions				
14. Managing Storage	Complete Chapter Lab				
15. Managing Advanced Storage	Read Foundation Topics				
15. Managing Advanced Storage	Complete Chapter Exercises				
15. Managing Advanced Storage	Review Key Topics				
15. Managing Advanced Storage	Complete Memory Tables				
15. Managing Advanced Storage	Define Key Terms				
15. Managing Advanced Storage	Answer Chapter Review Questions				
15. Managing Advanced Storage	Complete Chapter Lab				
16. Basic Kernel Management	Read Foundation Topics				
16. Basic Kernel Management	Complete Chapter Exercises				
16. Basic Kernel Management	Review Key Topics				
16. Basic Kernel Management	Complete Memory Tables				
16. Basic Kernel Management	Define Key Terms				
16. Basic Kernel Management	Answer Chapter Review Questions				
16. Basic Kernel Management	Complete Chapter Lab				
17. Managing and Understanding the Boot Procedure	Read Foundation Topics				
17. Managing and Understanding the Boot Procedure	Complete Chapter Exercises				
17. Managing and Understanding the Boot Procedure	Review Key Topics				
17. Managing and Understanding the Boot Procedure	Define Key Terms				
17. Managing and Understanding the Boot Procedure	Answer Chapter Review Questions				
17. Managing and Understanding the Boot Procedure	Complete Chapter Lab				
18. Essential Troubleshooting Skills	Read Foundation Topics				
18. Essential Troubleshooting Skills	Complete Chapter Exercises				
18. Essential Troubleshooting Skills	Review Key Topics				
18. Essential Troubleshooting Skills	Complete Memory Tables				
18. Essential Troubleshooting Skills	Define Key Terms				
18. Essential Troubleshooting Skills	Answer Chapter Review Questions				
18. Essential Troubleshooting Skills	Complete Chapter Lab				

19. An Introduction to Bash Shell Scripting	Read Foundation Topics				
19. An Introduction to Bash Shell Scripting	Complete Chapter Exercises				
19. An Introduction to Bash Shell Scripting	Review Key Topics				
19. An Introduction to Bash Shell Scripting	Define Key Terms				
19. An Introduction to Bash Shell Scripting	Answer Chapter Review Questions				
19. An Introduction to Bash Shell Scripting	Complete Chapter Lab				
20. Configuring SSH	Read Foundation Topics				
20. Configuring SSH	Complete Chapter Exercises				
20. Configuring SSH	Review Key Topics				
20. Configuring SSH	Complete Memory Tables				
20. Configuring SSH	Define Key Terms				
20. Configuring SSH	Answer Chapter Review Questions				
20. Configuring SSH	Complete Chapter Lab				
21. Managing Apache HTTP Services	Read Foundation Topics				
21. Managing Apache HTTP Services	Complete Chapter Exercises				
21. Managing Apache HTTP Services	Review Key Topics				
21. Managing Apache HTTP Services	Define Key Terms				
21. Managing Apache HTTP Services	Answer Chapter Review Questions				
21. Managing Apache HTTP Services	Complete Chapter Lab				
22. Managing SELinux	Read Foundation Topics				
22. Managing SELinux	Complete Chapter Exercises				
22. Managing SELinux	Review Key Topics				
22. Managing SELinux	Complete Memory Tables				
22. Managing SELinux	Define Key Terms				
22. Managing SELinux	Answer Chapter Review Questions				
22. Managing SELinux	Complete Chapter Lab				
23. Configuring a Firewall	Read Foundation Topics				
23. Configuring a Firewall	Complete Chapter Exercises				
23. Configuring a Firewall	Review Key Topics				
23. Configuring a Firewall	Complete Memory Tables				
23. Configuring a Firewall	Define Key Terms				
23. Configuring a Firewall	Answer Chapter Review Questions				
23. Configuring a Firewall	Complete Chapter Lab				

24. Accessing Network Storage	Read Foundation Topics				
24. Accessing Network Storage	Complete Chapter Exercises				
24. Accessing Network Storage	Review Key Topics				
24. Accessing Network Storage	Define Key Terms				
24. Accessing Network Storage	Answer Chapter Review Questions				
24. Accessing Network Storage	Complete Chapter Lab				
25. Configuring Time Services	Read Foundation Topics				
25. Configuring Time Services	Complete Chapter Exercises				
25. Configuring Time Services	Review Key Topics				
25. Configuring Time Services	Complete Memory Tables				
25. Configuring Time Services	Define Key Terms				
25. Configuring Time Services	Answer Chapter Review Questions				
25. Configuring Time Services	Complete Chapter Lab				
26. Final Preparation	Read Chapter				
26. Final Preparation	Review all Key Topics in all chapters				
26. Final Preparation	Complete all memory tables from Appendix B				
26. Final Preparation	Complete all exercises in all chapters				
26. Final Preparation	Complete all Labs in all chapters				
27. Theoretical Pre-Assessment Exam	Take the Pre-Assessment Exam found in the chapter				
	Take Practice Exam A found in the book				
	Take Practice Exam B found in the book				
	Take Practice Exam C found on the companion page				
	Take Practice Exam D found on the companion page				

Glossary

\$PATH A variable that contains a list of directories that are searched for executable files when a user enters a command.

- . The current directory. Its value can be requested using the **pwd** command.

A

absolute filename A filename that is complete and starts with the name of the root directory, including all directories up to the current file or directory.

access control list (ACL) In Linux permissions, a system that makes it possible to grant permissions to more than one user and more than one group. Access control lists also allow administrators to set default permissions for specific directories.

anacron A service that ensures that vital cron jobs can be executed when the server is down at the moment that the job normally should be executed. Can be considered an extension to cron.

AND A logical construction that can be used in scripts. In an AND construction, the second command is executed only after successful execution of the first command.

Application profile A collection of packages that may be used to install a specific version of software, according to a specific installation profile.

Application Stream A specific version of a yum module that can be installed as such.

archiving A system that ensures that data can be properly backed up.

at A service that can be used to schedule future jobs for one-time execution.

attribute A property that can be set to a file or directory and that will be enforced no matter which user with access permission accesses the file. For instance, a file that has the immutable (**i**) attribute set cannot be deleted, not even by the root user. However, the root user does have the capability to change the attribute, which would allow the root user to delete the file anyway.

audit log The main log file in `/var/log/audit/audit.log`, which by default contains all messages that are logged by the `auditd` service.

auditd A service that runs by default on Red Hat Enterprise Linux and can be configured to log very detailed information about what is happening on RHEL. Auditing is complementary to system logging and can be used for compliancy reasons. On RHEL, the auditing system takes care of logging SELinux-related messages, which makes it a relatively important system.

autofs A service that takes care of automatically mounting file systems at the moment that a specific directory is accessed. This service is very useful to ensure the automatic mounting of home directories for users in a centralized user management system, as can be implemented by the LDAP service.

automount The process that is started by the `autofs` service. *See* `autofs` for more details.

B

background process A process that is running on a system without actively occupying a console. Processes can be started in the background by adding a `&` after the command that starts the process. *See also* foreground process.

backup A copy of important data, which can be restored if at any point in time the original data gets lost.

Bash The default shell that is used on Red Hat Enterprise Linux.

Basic Input Output System (BIOS) The first software that is started when a computer starts on older IBM-compatible computers. Settings in the BIOS can be changed by using the BIOS setup program. *See also* Unified Extensible Firmware Interface (UEFI).

binary A numbering scheme that is based on bit values that can be on or off. Binary numbers are 0 and 1. Because binary numbers are difficult to use, decimal, hexadecimal, or octal numbers often are used.

BIOS *See* Basic Input Output System.

boot loader Program that is started as the very first thing while starting a computer and that takes care of loading the operating system kernel and `initramfs`.

Btrfs A general-purpose Linux file system that is expected to become the default file system on Red Hat Enterprise Linux in a future release.

C

cache In memory management, the area of memory where recently used files are stored. Cache is an important mechanism to speed up reads on servers.

capability A specific task that can be performed on Linux. User root has access to all capabilities; normal users have access to limited sets of capabilities only.

CentOS A Linux distribution that uses all Red Hat packages but has removed the Red Hat logo from all these packages to make it possible to distribute the software for free. CentOS is the best option for practicing for the RHCSA exam if you do not have access to RHEL.

certificate In PKI cryptography, contains the public key of the issuer of the certificate. This public key is signed with the certificate of a certificate authority, which guarantees its reliability.

certificate authority (CA) A commonly known organization that can be used to guarantee the reliability of PKI certificates. The certificate authority provides a certificate that can be used to sign public key certificates. Instead of using commonly known organizations, self-signed certificates can be used for internal purposes as well.

chrony The service that offers time synchronization services in Red Hat Enterprise Linux.

chroot An environment where a part of the file system is presented as if it were the root of the file system. Chroot is used as a security feature that hides part of the operating system that is not required by specific services.

CIFS *See* Common Internet File System.

cloud A computing platform that allows for flexible usage of hosted computing resources.

Common Internet File System (CIFS) The standardized version of the Microsoft Server Message Block (SMB) protocol, which is used to provide access to shared printers, files, and directories in a way that is compatible with Windows servers and clients. CIFS has become the de facto standard for file sharing in IT.

compression A technology that is used to reduce the size of files by analyzing redundant patterns and storing them more efficiently.

conditional loop In shell scripting, a set of commands that is executed only if a specific condition has been met.

connection (in network card configuration) A set of network configuration parameters that is associated to a network interface.

connection (in network communication) A session between two parties that has been initialized and will exist until the moment that the connection is tiered down.

console In Linux, the primary terminal where a user works. It is also a specific device with the name `/dev/console`.

context In SELinux, a label that is used to define the security attributes of users, processes, ports, and directories. These contexts are used in the SELinux policy to define security rules.

context switch When the CPU switches from executing one task to executing another task.

context type In SELinux, a label that identifies the SELinux properties of users, processes, ports, and processes.

Coordinated Universal Time (UTC) A time standard that is globally the same, no matter which specific time zone a user is in. UTC roughly corresponds to Greenwich Mean Time (GMT).

credentials file A file that can be used to mount CIFS file systems automatically from the `/etc/fstab` file. The credentials file is stored in a secure place, like the home directory of user root, and contains the username and password that are used to mount the remote file system.

cron A service that takes care of starting services repeatedly at specific times.

cryptography A technique used to protect data, often by converting information to an unreadable state, where keys are used to decipher the scrambled data. Cryptography is used not only to protect files while in transit but also to secure the authentication procedure.

D

deduplication A storage technology that analyzes data to be stored on disk and takes out duplicate patterns to allow for more efficient storage. Used in VDO. *See also* Virtual Data Optimizer (VDO).

default route The route that is used by default to forward IP packets to that have a destination on an external network.

dependency Generally, a situation where one item needs another item. Dependencies occur on multiple levels in Linux. In RPM package management, a dependency is a software package that needs to be present for another package to be installed. In systemd, a dependency is a systemd unit that must be loaded before another unit can be loaded.

dependency hell Situation where for package installation, other packages are needed, which by themselves could require dependencies as well. The problem of dependency hell has been fixed by the introduction of repository-based systems.

destination In rsyslog, specifies where log messages should be sent by the logging system. Destinations are often files, but can also be input modules, output modules, users, or hosts.

device A peripheral that is attached to a computer to perform a specific task.

device file A file that is created in the /dev directory and that is used to represent and interact with a device.

device mapper A service that is used by the Linux kernel to communicate with storage devices. Device mapper is used by LVM, multipath, and other devices, but not by regular hard disks. Device files that are created by device mapper can be found in the /dev/mapper directory.

directory A folder in the file system that can be used to store files in an organized manner.

disabled mode The SELinux mode in which SELinux is completely deactivated.

distribution A Linux version that comes with its own installation program or which is ready for usage. Because Linux is a collection of different tools and other components, the Linux distribution gathers these tools and other components, may or may not enhance them, and distributes them so that users do not have to gather all the different components for themselves.

dmesg Utility that can be used to read the kernel ring buffer, which contains log messages that were generated by the Linux kernel.

dracut A utility that is used to generate the initramfs, an essential part of the Linux operating system that contains drivers and other vital files required to start a Linux system.

dynamic route A network route that is managed by an automatic routing protocol.

E

enforcing mode The SELinux mode where SELinux is fully operational and applies all restrictions that have been configured for a specific system.

environment The collection of settings that users or processes are using to do their work.

epoch time In Linux, the number of seconds that have passed since epoch (corresponds to midnight on January 1, 1970). Some utilities write epoch time instead of real clock time.

escaping In a shell environment, using special syntax to ensure that specific characters are not interpreted by the shell. Escaping may be necessary to show specific characters onscreen or to ensure that regular expression metacharacters are not interpreted by the bash shell first.

export In NFS, refers to a directory that is shared on an NFS server to allow access to other servers.

Ext2, 3, and 4 Three different versions of the Ext file system. Up to RHEL 6, Ext4 was the default file system. It is now considered inadequate for modern storage needs, which is why Ext4 in RHEL 7 has been replaced by XFS as the default file system.

extended partition A solution to create more than four partitions on an MBR disk. On MBR disks, a maximum of four partitions can be stored in the partition table. To make it possible to go beyond that number, one of the four partitions can be created as an extended partition. Within an extended partition, logical partitions can be created, which will perform just like regular partitions, allowing system administrators to create more partitions.

external command A command that exists as a file on disk.

F

facility In rsyslogd, the source where log information comes from. A strictly limited number of facilities have been defined in rsyslogd.

Fedora The free and open source Linux distribution that is sponsored by Red Hat. In Fedora, new features are provided and tested. Some of these features will be included in later releases of Red Hat Enterprise Linux.

FHS *See* Filesystem Hierarchy Standard.

file descriptor A pointer that is used by a Linux process to refer to files that are in use by the process.

file system A logical structure that is created on a storage device. In a Linux file system, inodes are used for file system administration, and the actual data is written to blocks. *See also* inode.

Filesystem Hierarchy Standard (FHS) A standard that defines which Linux directories should be used for which purpose. Read man 7 hier for a specification of the FSH.

firewall A solution that can be used to filter packets on a network. Firewalls are used to ensure that only authorized traffic can reach a system. A firewall can be offered through the Linux kernel Netfilter functionality but often is also offered as an appliance on the network.

firewalld The modern service (replacing iptables) that is used in RHEL 7 and RHEL 8 to implement firewalling based on the Linux kernel firewalling framework.

folder Also referred to as a directory, a structure in the file system used to organize files that belong together.

foreground process Linux processes that are started by users can be started in the foreground or in the background. If a process has been started as a foreground process, no other processes can be started in the same terminal until it finishes or is moved to the background. *See also* background process.

fstab A configuration file that is used on Linux to mount file systems automatically when the system starts.

G

GECOS A field in the `/etc/passwd` file that can be used to store personal data about a user on the Linux operating system. GECOS originally stood for General Electric Comprehensive Operating Supervisor.

global unique ID (GUID) An identification number that consists of parts that ensure that it is globally unique.

GPT *See* GUID Partition Table.

group A collection of items. In user management, groups are used to assign permissions to multiple users simultaneously. In Linux, every user is a member of at least one group.

group owner The group that has been set as the owner of a file or a directory. On Linux, every file and directory has a user owner and a group owner. Group ownership is set when files are created, and unless configured otherwise, it is set to the primary group of the user who creates the file.

GRUB *See* GRUB 2.

GRUB 2 The boot loader that is installed on most systems that need to start Linux. GRUB 2 provides a boot prompt from which different kernel boot options can be entered, which is useful if you need to troubleshoot the boot procedure.

GUID *See* global unique ID.

GUID Partition Table (GPT) A modern solution to store partitions on a hard disk, as opposed to the older MBR partition table. In GUID partitions, a total of 128 partitions can be created, and no difference exists between primary, extended, and logical partitions anymore.

gzip One of the most common utilities that is used for compression and decompression of files on Linux.

H

hard link A name associated with an inode. Inodes are used to store Linux files. An inode contains the complete administration of the file, including the blocks in which the file is stored. A file that does not have at least one hard link is considered a deleted file. To increase file accessibility, more than one hard link can be created for an inode.

hardware time The time that is provided by computer hardware, typically the BIOS clock. When a Linux system boots, it sets the software time based on the hardware time. Because hardware time often is inaccurate, most Linux systems use the Network Time Protocol (NTP) to synchronize the system time with a reliable time source.

hexadecimal A 16-based numbering system that is based on groups of 4 bytes. Hexadecimal numbers start with the range 0 through 9, followed by A through F. Because hexadecimal is much more efficient in computer technology, hexadecimal numbers are often used. In IPv6, IP addresses are written as hexadecimal numbers.

hypervisor A piece of computer software, firmware, or hardware that creates and runs virtual machines. In Linux, KVM is used as the common hypervisor software.

I

inheritance In permission management, refers to the situation where new files that are created in a directory inherit the permission settings from the parent directory.

init The first process that is started once the Linux kernel and initramfs have been loaded. From the init process, all other processes are started. As of RHEL 7, the init process has been replaced by systemd.

initramfs The initial RAM file system. Contains drivers and other files that are needed in the first stages of booting a Linux system. On Red Hat Enterprise Linux, the initramfs is generated during installation and can be manually re-created using the **dracut** utility.

inode Contains the complete administration of a file. Every Linux file has an inode, and the inode contains all properties of the file but not the filename.

input module In rsyslog, a module that allows rsyslog to receive log messages from specific sources.

interface In Linux networking, the set of configuration parameters that can be activated for a specific device. Several interface configurations can exist for a device, but only one interface can be active at a time for a device.

internal command A command that is a part of the shell and does not exist as a file on disk.

Internet Protocol (IP) The primary communications protocol that is used by computers for communication. The Internet Protocol exists in two versions (IPv4 and IPv6). Apart from node addressing, it defines routing, which enables nodes to contact one another.

IP *See* Internet Protocol.

iptables An older solution to create firewall rules on the Linux operating system. Interfaces with the Netfilter Linux kernel firewalling functionality and was the default solution to create software firewalls on earlier versions of RHEL. As of RHEL 7, iptables has been replaced by firewalld.

IPv4 Version 4 of the Internet protocol. It was developed in the 1970s and introduced in 1981. It allows a theoretical maximum of about 4 billion nodes to be addressed by using a 32-bit address space. It is still the most important IP version in use.

IPv6 Version 6 of the Internet protocol. It was developed in the 1990s to address the shortage in IPv6 addresses. It uses a 128-bit address space that allows for addressing 3,4e38 nodes and thus is considered a virtually unlimited address space.

iteration In shell scripting, one time of many that a conditional loop has been processed until the desired result has been reached.

J

job In a Linux shell, a task running in the current terminal. Jobs can be started in the foreground and in the background. Every job is also visible as a process.

journalctl The part of systemd that takes care of logging messages.

journald The part of systemd that takes care of logging information about events that have been happening. The introduction of journald ensures that information about all services can be logged, regardless of how the service itself is configured to deal with information that is to be logged.

K

kernel The central component of the operating system. It manages I/O requests from software and translates them into data processing instructions for the hardware in the computer.

kernel ring buffer A part of memory where messages that are generated by the kernel are stored. The **dmesg** command enables you to read the contents of the kernel ring buffer.

kernel space The part of memory that is reserved for running privileged instructions. Kernel space is typically accessible by the operating system kernel, kernel extensions, and most device drivers. Applications normally run in user space, which ensures that a faulty application cannot crash the computer system.

Kernel-based Virtual Machine (KVM) The Linux kernel module that acts as a hypervisor and makes it possible to run virtual machines directly on top of the Linux kernel.

key-based login In SSH, login that uses public/private keys to proof the identity of the user who wants to log in. Key-based login is generally considered more secure than password-based login.

kill A command that can be used to send a signal to a Linux process. Many signals are defined (see `man 7 signal`), but only a few are commonly used, including `SIGTERM` and `SIGKILL` that both are used to stop processes.

KVM See Kernel-based Virtual Machine.

L

label A name that can be assigned to a file system. Using labels can be a good idea, because once a label is assigned, it will never be changed, which guarantees that the file system can still be mounted, even if other parameters such as the device name have changed. However, UUIDs are considered safer than labels because the chance of having a duplicate label by accident is much higher than the chance of having a duplicate UUID. See *also* universally unique ID (UUID).

line anchor In regular expressions, a character that refers to a specific position in a line.

Linux A UNIX-like operating system that consists of a kernel that was originally developed by Linus Torvalds (hence the name Linux). A current Linux operating system consists of a kernel and lots of open source tools that provide a complete operating system. Linux is packaged in the form of a distribution. Currently, Red Hat Enterprise Linux is among the most widely used Linux distributions.

log rotation A service that ensures that log files cannot grow too big. Log files are monitored according to specific parameters, such as a maximum age or size.

Once this parameter is reached, the log file will be closed and a new log file will be opened. Old log files are kept for a limited period and will be removed, often after only a couple of weeks.

logical extent The building block that is used in LVM to create logical volumes. It normally has a size of a few megabytes that corresponds to the size of the physical extents that are used.

logical partition A partition that is created in an extended partition. *See also* extended partition.

logical volume In LVM, the entity on which a file system is created. Logical volumes are often used on RHEL because they offer important advantages, such as the option to dynamically resize the logical volume and the file system that it hosts.

Logical Volume Manager (LVM) The software that makes it possible to work with logical volumes.

login shell The shell that is opened directly after a user has logged in.

LVM *See* Logical Volume Manager.

M

masquerading A solution that enables a private IP address range that is not directly accessible from outside networks to be accessed by using one public IP address that is exposed on a router.

Master Boot Record (MBR) On a BIOS system, the first 512 bytes on the primary hard disk. It contains a boot loader and a partition table that give access to the different partitions on the hard disk of that computer.

MBR *See* Master Boot Record.

module A piece of snap-in code. Modules are used by several systems on Linux, such as the kernel, GRUB 2, rsyslog, and more. Via modules, Linux components can be extended easily, and adding functionality does not require a total rewrite of the software.

module (in YUM) A collection of software packages that can be managed as one entity and can contain different versions of a software solution.

mount A connection that is made between a device and a directory. To access files on specific storage devices, the storage device needs to be mounted on a directory. This sets up the specified directory as the access point to files on the storage device. Mounts are typically organized by the systems administrator and are not visible to end users.

multiplier In regular expressions, a character that indicates that multiple of the previous character are referred to.

N

netfilter The part of the Linux kernel that implements firewalling.

netmask *See* subnet mask.

Network Address Translation (NAT) *See* masquerading.

Network File System (NFS) A common UNIX solution to export physical file systems to other hosts on the network. The other hosts can mount the exported NFS directory in their local file system.

network time Time that is provided on the network.

Network Time Protocol (NTP) A standard that is used to provide reliable time to servers in a network. NTP on RHEL 8 is implemented by the chronyd service.

NFS *See* Network File System.

nftables The service that manages kernel firewalling. It is a replacement of the older iptables service.

nice A method to change the priority of Linux processes. A negative nice value will make the process more aggressive, giving it a higher priority (which is expressed by a lower priority number); a positive nice value will make a process less eager so that it gives priority to other processes.

NTP *See* Network Time Protocol.

O

octal A numbering scheme that uses the numbers 0 through 7 only. Used when working with Linux permissions using the **umask** or the **chmod** commands.

OR A logical operation where the second command is executed only if the first command is not able to execute.

output module In rsyslog, a module that is used to send log messages to a specific destination. Output modules make rsyslogd flexible and allow for the usage of log destinations that are not native to rsyslog.

ownership In file system permissions, the basis of the effective permissions that a user has. Every file has a user owner and a group owner assigned to it.

P

package A bundle that is used to distribute software. A package typically contains a compressed archive of files and metadata that includes instructions on how to install those files.

package group A group of packages that can be installed as such using the **yum groups install** command.

package group (in yum) A group of software packages that can be installed with a single command.

pager A program that can be used to browse page by page through a text file. The **less** utility provides one of the most common Linux pagers.

parent shell The environment from which a shell script or program is started. Processes or child scripts will inherit settings from the parent shell.

partition A subdivision of a hard disk on which a file system can be created to mount it into the directory structure.

passphrase Basically a password, but is supposed to be longer and more secure than a password.

password A token that is used in authentication. The password is a secret word that can be set by individual users and will be stored in an encrypted way.

path The complete reference to the location of a file.

permissions Attributes that can be set on files or directories to allow users or groups access to these files or directories.

permissive mode The SELinux mode where nothing is blocked but everything is logged in the audit log. This mode is typically used for troubleshooting SELinux issues.

physical extent The physical building block that is used when creating LVM physical volumes. Typically, the size is multiple megabytes.

physical volume The foundation building block of an LVM configuration. The physical volume typically corresponds to a partition or a complete disk device.

PID *See* process identification number.

pipe A structure that can be used to forward the output of one command to be used as input for another command.

policy *See* SELinux policy.

port A number that is used by a process to offer access to the process through a network connection.

port forwarding A firewalling technique where traffic that is coming in on a specific port is forwarded to another port that may be on the same host or on a different host.

Portable Operating System Interface (POSIX) A standard that was created to maintain compatibility between operating systems. The standard mainly applies to UNIX and guarantees that different flavors of Linux and UNIX are compatible with one another.

portmapper A remote procedure call service that needs to run on systems that provide RPC services. Portmapper uses dynamic ports that do not correspond to specific TCP or UDP ports; the service will pick a UDP or TCP port that will be used as long as the process is active. When restarted, chances are that different ports are used. They need to be mapped to fixed UDP and TCP ports in order to make it possible to open the firewall for these ports. Portmapper is still used by components of the NFS service.

POSIX *See* Portable Operating System Interface.

primary group The group that is listed in the group membership field for a user in `/etc/passwd`. Every Linux user is a member of a primary group. Apart from that, users can be made a member of secondary groups as well.

primary partition In MBR, one of a maximum of four partitions that can be created in the Master Boot Record. *See* also extended partition.

priority (in process handling) Specifies the importance of a process. Process priority is expressed with a number (which can be modified using **nice**). Processes with a lower priority number are serviced before processes with a higher priority number.

priority (in rsyslog) Used to specify the severity of a logged event. Based on the severity, specific actions can be taken.

private key In public/private key encryption, the key that is used to generate encrypted data.

privileged user *See* root.

proc A kernel interface that provides access to kernel information and kernel tunables. This interface is available through the `/proc` file system.

process A task that is running on a Linux machine. Roughly, a process corresponds to a program, although one program can start multiple processes.

process identification number (PID) A unique number that is used to identify a process running on a Linux system.

profile In tuned, a collection of performance settings that can easily be applied.

protocol A set of rules that is used in computing, such as in computer networking to establish communications between two computers.

Pseudo Root File System In the current NFS release NFSv4, a solution where multiple shares are exported by an NFS server. Instead of mounting each individual share, the NFS client mounts the root file system on the NFS server, which gives access to all shares the client is entitled to.

pseudo root mount In NFSv4, a mount of the root directory system or another high-level directory that hasn't been specifically exported by the NFS server, but which gives access to all exported file systems the client has access to.

public key In cryptography, the key that is typically sent by a server to a client so that the client can send back encrypted data.

PV See physical volume.

Q

queue In process management, where processes wait before they can be executed.

R

real-time clock The hardware clock that is installed on the computer motherboard.

reboot The procedure of stopping the computer and starting it again.

Red Hat Customer Portal The platform that Red Hat offers to provide patches for customers that have an active subscription. To provide these patches and updates, Red Hat Network provides the repositories that are needed for this purpose.

Red Hat Enterprise Linux (RHEL) The name of the software that Red Hat sells subscriptions for. It is available in a server edition and a desktop edition.

Red Hat Package Manager The name for the package format that is used on RHEL for software packages and for the Package Management software. RPM has become the standard for package management on many other Linux distributions as well.

reference clock A clock that is used as a time source in an NTP time configuration. Typically, a reference clock is a highly reliable clock on the Internet, but it can be an internal clock on the computer's motherboard as well.

regular expression A search pattern that allows users to search text patterns in a flexible way. Not to be confused with shell metacharacters.

relative filename A filename that is relative to a directory that is not the root directory.

Remote Procedure Calls (RPC) A method for interprocess communication that allows a program to execute code in another address space. Remote Procedure Calls is an old protocol and as such is still used in the Network File System.

repository An installation source that contains installable packages and an index that contains information about the installable packages so that the installation program yum can compare the version of packages currently installed with the version of packages available in the repository.

resident memory Memory pages that are in use by a program.

resolver The DNS client part that contains a list of DNS servers to contact to resolve DNS queries.

RHEL *See* Red Hat Enterprise Linux.

rich rules Rules in firewalld that allow the usage of a more complicated syntax so that more complex rules can be defined.

root The privileged user account that is used for system administration tasks. User root has access to all capabilities, which means that permissions do not apply to user root and the root user account is virtually unlimited.

root directory The starting point of the file system hierarchy, noted as /.

RPC *See* Remote Procedure Calls.

RPM *See* Red Hat Package Manager.

RTC *See* real-time clock.

rsyslogd The generic daemon that logs messages.

S

Samba The name for the Linux service that implements the SMB protocol.

SAN *See* storage-area network.

scheduler The part of the Linux kernel that monitors the queue of runnable processes and allocates CPU time to these processes.

Scientific Linux A Linux distribution that is based on the Red Hat packages from which the Red Hat logo has been removed. A very good choice for people that are looking for a freely available alternative to Red Hat Enterprise Linux.

secondary group A group that a user is a member of but which membership is not defined in the /etc/passwd file. When creating new files, the secondary group will not automatically become the owner of those files.

Secure Shell (SSH) A solution that allows users to open a shell on a remote server where security is implemented by using public/private key cryptography.

Secure Sockets Layer (SSL) *See* Transport Layer Security (TLS).

SELinux A Linux kernel security module that provides a mechanism for supporting access control security policies.

SELinux Policy The collection of rules that is used to define SELinux security.

Server Message Block (SMB) An application-level protocol that is used to provide shared access to files, printers, and serial ports, which on Linux is implemented in the Samba server.

services (in firewalld) A configuration of firewall settings that is used to allow access to specific processes.

services (in systemd) Processes that need to be started to provide specific functionality.

share A directory to which remote access is configured using a remote file system protocol such as NFS or CIFS.

shebang Used in a script to indicate which shell should be used for executing the code in the shell script. If no shebang is used, the script code will be interpreted by the parent shell, which may lead to errors in some cases. A shebang starts with a #, which is followed by a ! and the complete pathname of the shell, such as #!/bin/bash.

shell The environment from which commands can be executed. Bash is the default shell on Linux, but other shells exist as well.

shell metacharacters Characters such as *, ?, and [a-z] that allow users to refer to characters in filenames in a flexible way.

signal An instruction that can be sent to a process. Common signals exist, such as SIGTERM and SIGKILL, but the Linux kernel allows a total of 32 different signals to be used. To send a signal to a process, use the **kill** command.

SMB *See* Server Message Block.

snapshot A “photo” of the actual state of a file system.

software time *See* system time.

source context In SELinux, the context of the processes or users that initiate an action. A context in SELinux is a label that identifies allowed operations. Everything in an SELinux environment has a context.

SSH *See* Secure Shell.

standard error (STDERR) The default location where a program sends error messages.

standard input (STDIN) The default location where a program gets its input.

standard output (STDOUT) The default location where a program sends its regular output.

static route A route that is defined manually by a network administrator.

STDERR *See* standard error.

STDIN *See* standard input.

STDOUT *See* standard output.

storage-area network (SAN) A solution where disk devices are shared at a block level over the network. As such, they can be used in the same way as local disk devices on a Linux system. iSCSI and Fibre Channel are the common SAN protocols.

Stratis The new volume managing file system in RHEL 8.

stratum In time synchronization, used to indicate the distance between a server and an authoritative Internet time source.

subnet mask A logical subdivision of an IP network.

subshell A shell that is started from another shell. Typically, by running a shell script a subshell is started.

symbolic link A special type of file that contains a reference to another file or directory in the form of an absolute or relative path.

sysfs The kernel interface that is mounted on the /sys directory and which is used to provide access to parameters that can be used for managing hardware settings.

system time The time that is maintained by the operating system. When a Linux system boots, system time is set to the current hardware time, and while the operating system is running, it is often synchronized using the Network Time Protocol (NTP).

systemd The service manager on RHEL 8. systemd is the very first process that starts after the kernel has loaded, and it takes care of starting all other processes and services on a Linux system.

T

tainted kernel A kernel in which unsupported kernel modules have been loaded.

tar The Tape Archiver; the default Linux utility that is used to create and extract backups.

target In systemd, a collection of unit files that can be managed together.

target context The SELinux context that is set to a target object, such as a port, file, or directory.

terminal Originally, the screen that was used by a user to type commands on. On modern Linux systems, pseudo terminals can be used as a replacement. A pseudo terminal offers a shell window from which users enter the commands that need to be executed.

thin allocation In storage, an approach that enables the system to present more storage to the storage user than what is really available by using smart technologies to store data, like deduplication.

thread A thread can be used as a subdivision of a process. Many processes are single threaded, which means that process is basically one entity that needs to be serviced. On a multicore or multi-CPU computer system, working with multithreaded processes makes sense. That way, the different cores can be used to handle the different threads, which allows a process to benefit from multicore or multithreaded environments.

time stamp An identifier that can be used on files, database records, and other types of data to identify when the last modification has been applied. Many services rely on time stamps. To ensure that time stamped-based systems work properly, time synchronization needs to be configured.

time synchronization A system that ensures that multiple servers are using the exact same time. To accomplish time synchronization, it is common to use an external time server, as defined in the Network Time Protocol (NTP).

timer A systemd unit type that can be used as an alternative to cron jobs and run units at a specific time.

TLS See Transport Layer Security.

Transport Layer Security (TLS) A cryptographic protocol that is created to ensure secured communications over a computer network. In TLS, public and private keys are used, and certificates authenticate the counterparty. TLS was formerly known as SSL.

TTY A program that provides a virtual terminal on Linux. Every terminal still has a TTY name, which is either tty1-6 for virtual TTYs or /dev/pts/0-nn for pseudo terminals.

tuned A service on RHEL that enables administrators to easily apply performance settings by using profiles.

U

udev A service that works with the Linux kernel to initialize hardware.

UEFI *See* Unified Extensible Firmware Interface.

umask An octal value that defines the default permissions as a shell property.

umount The command that is used to decouple a file system from the directory on which it is mounted.

Unified Extensible Firmware Interface (UEFI) A replacement of the Basic Input Output System used on older IBM-compatible computers as the first program that runs when the computer is started. UEFI is the layer between the operating system and the computer firmware.

unit In systemd, refers to an item that is managed by systemd. Different types of units exist, including service, path, mount, and target units.

universally unique ID (UUID) An identification number that consists of a long random hexadecimal number and which is globally unique.

unprivileged user A regular non-root user account to which access restrictions apply, as applied by permissions.

user An entity that is used on Linux to provide access to specific system resources. Users can be used to represent people, but many services also have a dedicated user account, which allows the service to run with the specific permissions that are needed for that service.

user space The area of memory that is accessible by application software that has been started with non-root privileges.

UTC *See* Coordinated Universal Time.

UUID *See* universally unique ID.

V

value The data that is assigned to a specific property, variable, or record.

variable A label that contains a specific value that can be changed dynamically. In scripting, variables are frequently used to allow the script to be flexible.

VDO See Virtual Data Optimizer.

VFAT The Linux kernel driver that is used to access FAT-based file systems. FAT is a commonly used file system in Windows environments. The Linux VFAT driver allows usage of this file system.

VG See volume group.

Virtual Data Optimizer (VDO) A new advanced storage solution in RHEL 8 that compresses data by using deduplication and can therefore be used for thin allocation of storage volumes.

virtual host In the Apache web server, a collection of configuration settings that is used to address a web server. What makes it a virtual host is that one installation of the Apache web server can be configured with multiple virtual hosts, which allows administrators to run multiple websites on one Apache server.

virtual memory The total amount of addressable memory. Virtual memory is called *virtual* memory because it does not refer to memory that really exists. Its only purpose is to make sure that Linux programs can set an address pointer that is unique and not in use by other programs.

volume group The abstraction layer that in Logical Volume Manager is used to represent all available storage presented by physical volumes from which logical volumes can be created.

W

want An indication for a systemd unit file that it is supposed to be started from a specific systemd target.

wildcard The * character, which in a shell environment refers to an unlimited number of any characters.h.pakaein

X

XFS A high-performance 64-bit file system that was created in 1993 by SGI and which in RHEL 8 is used as the default file system.

Y

Yellowdog Update, Modified The full name for YUM, the meta package handler that on RHEL 8 is used to install packages from yum repositories.

Yum *See* Yellowdog Update, Modified.

Z

zombie A process that has lost contact with its parent and for that reason cannot be managed using regular tools.

zone In firewalld, a collection of one or more network interfaces that specific firewalld rules are associated with.