



Red Hat Training and Certification

Student Workbook (ROLE)

Red Hat Enterprise Linux 9.3 RH134

Red Hat System Administration II

Edition 2



Red Hat
Learning Community

Join a community dedicated to learning open source

The Red Hat® Learning Community is a collaborative platform for users to accelerate open source skill adoption while working with Red Hat products and experts.



Network with tens of thousands of community members




Engage in thousands of active conversations and posts



Join and interact with hundreds of certified training instructors



Unlock badges as you participate and accomplish new goals



This knowledge-sharing platform creates a space where learners can connect, ask questions, and collaborate with other open source practitioners.

Access free Red Hat training videos

Discover the latest Red Hat Training and Certification news

Connect with your instructor – and your classmates – before, after, and during your training course.

Join peers as you explore Red Hat products

Join the conversation learn.redhat.com



Copyright © 2020 Red Hat, Inc. Red Hat, Red Hat Enterprise Linux, the Red Hat logo, and Ansible are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

Red Hat System Administration II

Red Hat Enterprise Linux 9.3 RH134

Red Hat System Administration II

Edition 2 20241029

Publication date 20241029

Authors: Ashish Lingayat, Alejandra Ramírez Palacios, Alex Callejas,
Antonio Marí Romero, Bernardo Gargallo, Dallas Spohn, Ed Parenti,
Jacob Pelchat, Mike Kelly, Morgan Weetman, Patrick Gomez,
Shreya Dchange

Course Architects: Philip Sweany, Steven Bonneville

DevOps Engineers: Artur Glogowski, Trey Feagle

Editors: Julian Cable, Rachel Lee

© 2024 Red Hat, Inc.

The contents of this course and all its modules and related materials, including handouts to audience members, are © 2024 Red Hat, Inc.

No part of this publication may be stored in a retrieval system, transmitted or reproduced in any way, including, but not limited to, photocopy, photograph, magnetic, electronic or other record, without the prior written permission of Red Hat, Inc.

This instructional program, including all material provided herein, is supplied without any guarantees from Red Hat, Inc. Red Hat, Inc. assumes no liability for damages or legal action arising from the use or misuse of contents or details contained herein.

If you believe Red Hat training materials are being used, copied, or otherwise improperly distributed, please send email to training@redhat.com [mailto:training@redhat.com] or phone toll-free (USA) +1 (866) 626-2994 or +1 (919) 754-3700.

Red Hat, Red Hat Enterprise Linux, the Red Hat logo, JBoss, OpenShift, Fedora, Hibernate, Ansible, RHCA, RHCE, RHCSA, Ceph, and Gluster are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

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

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

XFS® is a registered trademark of Hewlett Packard Enterprise Development LP or its subsidiaries in the United States and/or other countries.

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

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

The OpenStack word mark and the Square O Design, together or apart, are trademarks or registered trademarks of OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. Red Hat, Inc. is not affiliated with, endorsed by, or sponsored by the OpenStack Foundation or the OpenStack community.

All other trademarks are the property of their respective owners.

Contributors: Adarsh Krishnan, David Sacco, Hemant Chauhan, Jonika Shank, Roberto Velazquez, Sajith Eyamkuzhy, Samik Sanyal, Shashi Singh, Yuvaraj Balaraju

Document Conventions	xi
Admonitions	xi
Inclusive Language	xii
Introduction	xiii
Red Hat System Administration II	xiii
Orientation to the Classroom Environment	xiv
Performing Lab Exercises	xviii
1. Improve Command-line Productivity	1
Write Simple Bash Scripts	2
Guided Exercise: Write Simple Bash Scripts	6
Loops and Conditional Constructs in Scripts	9
Guided Exercise: Loops and Conditional Constructs in Scripts	15
Match Text in Command Output with Regular Expressions	17
Guided Exercise: Match Text in Command Output with Regular Expressions	26
Lab: Improve Command-line Productivity	29
Summary	35
2. Schedule Future Tasks	37
Schedule a Deferred User Job	38
Guided Exercise: Schedule a Deferred User Job	40
Schedule Recurring User Jobs	43
Guided Exercise: Schedule Recurring User Jobs	46
Schedule Recurring System Jobs	49
Guided Exercise: Schedule Recurring System Jobs	52
Manage Temporary Files	55
Guided Exercise: Manage Temporary Files	58
Quiz: Schedule Future Tasks	61
Summary	65
3. Analyze and Store Logs	67
Describe System Log Architecture	68
Quiz: Describe System Log Architecture	70
Review Syslog Files	74
Guided Exercise: Review Syslog Files	79
Review System Journal Entries	81
Guided Exercise: Review System Journal Entries	87
Preserve the System Journal	90
Guided Exercise: Preserve the System Journal	93
Maintain Accurate Time	96
Guided Exercise: Maintain Accurate Time	100
Lab: Analyze and Store Logs	104
Summary	109
4. Archive and Transfer Files	111
Manage Compressed tar Archives	112
Guided Exercise: Manage Compressed tar Archives	117
Transfer Files Between Systems Securely	119
Guided Exercise: Transfer Files Between Systems Securely	122
Synchronize Files Between Systems Securely	125
Guided Exercise: Synchronize Files Between Systems Securely	128
Lab: Archive and Transfer Files	131
Summary	136
5. Tune System Performance	137
Adjust Tuning Profiles	138
Guided Exercise: Adjust Tuning Profiles	145

Influence Process Scheduling	150
Guided Exercise: Influence Process Scheduling	154
Lab: Tune System Performance	158
Summary	164
6. Manage SELinux Security	165
Change the SELinux Enforcement Mode	166
Guided Exercise: Change the SELinux Enforcement Mode	171
Control SELinux File Contexts	174
Guided Exercise: Control SELinux File Contexts	179
Adjust SELinux Policy with Booleans	182
Guided Exercise: Adjust SELinux Policy with Booleans	184
Investigate and Resolve SELinux Issues	187
Guided Exercise: Investigate and Resolve SELinux Issues	192
Lab: Manage SELinux Security	196
Summary	202
7. Manage Basic Storage	203
Add Partitions, File Systems, and Persistent Mounts	204
Guided Exercise: Add Partitions, File Systems, and Persistent Mounts	213
Manage Swap Space	217
Guided Exercise: Manage Swap Space	221
Lab: Manage Basic Storage	225
Summary	233
8. Manage Storage Stack	235
Create and Extend Logical Volumes	236
Guided Exercise: Create and Extend Logical Volumes	246
Manage Layered Storage	252
Guided Exercise: Manage Layered Storage	258
Lab: Manage Storage Stack	264
Summary	271
9. Access Network-Attached Storage	273
Manage Network-Attached Storage with NFS	274
Guided Exercise: Manage Network-Attached Storage with NFS	277
Automount Network-Attached Storage	280
Guided Exercise: Automount Network-Attached Storage	285
Lab: Access Network-Attached Storage	291
Summary	297
10. Control the Boot Process	299
Select the Boot Target	300
Guided Exercise: Select the Boot Target	306
Reset the Root Password	309
Guided Exercise: Reset the Root Password	313
Repair File-system Issues at Boot	315
Guided Exercise: Repair File-system Issues at Boot	318
Lab: Control the Boot Process	321
Summary	327
11. Manage Network Security	329
Manage Server Firewalls	330
Guided Exercise: Manage Server Firewalls	338
Control SELinux Port Labeling	342
Guided Exercise: Control SELinux Port Labeling	345
Lab: Manage Network Security	349
Summary	357

12. Install Red Hat Enterprise Linux	359
Install Red Hat Enterprise Linux	360
Guided Exercise: Install Red Hat Enterprise Linux	365
Automate Installation with Kickstart	368
Guided Exercise: Automate Installation with Kickstart	377
Install and Configure Virtual Machines	380
Quiz: Install and Configure Virtual Machines	385
Lab: Install Red Hat Enterprise Linux	387
Summary	393
13. Run Containers	395
Introduction to Containers	396
Quiz: Introduction to Containers	400
Creating Containers with Podman	404
Guided Exercise: Creating Containers with Podman	410
Container Image Registries	416
Guided Exercise: Container Image Registries	423
Managing the Container Lifecycle	427
Guided Exercise: Managing the Container Lifecycle	436
Lab: Run Containers	442
Summary	449
14. Comprehensive Review	451
Comprehensive Review	452
Lab: Fix Boot Issues and Maintain Servers	455
Lab: Configure and Manage File Systems and Storage	461
Lab: Configure and Manage Server Security	467
Lab: Run Containers	477

Document Conventions

This section describes various conventions and practices that are used throughout all Red Hat Training courses.

Admonitions

Red Hat Training courses use the following admonitions:



References

These describe where to find external documentation that is relevant to a subject.



Note

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



Important

Important sections provide details of information that is easily missed: configuration changes that apply only to the current session, or services that need restarting before an update applies. Ignoring these admonitions will not cause data loss, but might cause irritation and frustration.



Warning

Do not ignore warnings. Ignoring these admonitions will most likely cause data loss.

Inclusive Language

Red Hat Training is currently reviewing its use of language in various areas to help remove any potentially offensive terms. This is an ongoing process and requires alignment with the products and services that are covered in Red Hat Training courses. Red Hat appreciates your patience during this process.

Introduction

Red Hat System Administration II

This course is specifically designed for students who have completed the *Red Hat System Administration I (RH124)* course. The *Red Hat System Administration II (RH134)* course focuses on the key tasks that are needed to become a full-time Linux administrator and to validate those skills via the *Red Hat Certified System Administrator* exam. This course goes deeper into Enterprise Linux administration, including file systems, partitioning, logical volumes, SELinux, firewalls, and troubleshooting.

Course Objectives

- Expand on skills that were gained during the *Red Hat System Administration I (RH124)* course.
- Build skills that an RHCSA-certified Red Hat Enterprise Linux system administrator needs.

Audience

- This course is singularly designed for students who have completed *Red Hat System Administration I (RH124)*. Given the organization of topics, it is not appropriate for students to use RH134 as a curriculum entry point. Students who have not taken a previous Red Hat course are encouraged to take either *Red Hat System Administration I (RH124)* if they are new to Linux, or the *RHCSA Rapid Track (RH199)* course if they are experienced with Enterprise Linux administration.

Prerequisites

- To have taken the *Red Hat System Administration I (RH124)* course, or equivalent knowledge.

Orientation to the Classroom Environment




Figure 0.1: Classroom environment

In this course, the main computer system for hands-on learning activities is **workstation**. Students also use two other machines for these activities: **servera** and **serverb**. All three systems are in the `lab.example.com` DNS domain.

All student computer systems have a standard user account, `student`, which has the password `student`. The root password on all student systems is `redhat`.

Classroom Machines

Machine name	IP address	Role
<code>bastion.lab.example.com</code>	172.25.250.254	Gateway system to connect the student private network to the classroom server (must always be running)
<code>classroom.example.com</code>	172.25.254.254	Server that hosts the required classroom materials
<code>workstation.lab.example.com</code>	172.25.250.9	Graphical workstation for student use
<code>servera.lab.example.com</code>	172.25.250.10	Managed server "A"
<code>serverb.lab.example.com</code>	172.25.250.11	Managed server "B"

The primary function of **bastion** is to act as a router between the network that connects the student machines and the classroom network. If **bastion** is down, then other student machines can access only systems on the individual student network.

Several systems in the classroom provide supporting services. Two servers, `content.example.com` and `materials.example.com`, are sources for software and lab materials in hands-on activities. Information about how to use these servers is provided in the instructions for those activities. The `workstation` virtual machine provides these activities. Both `classroom` and `bastion` must always be running for proper use of the lab environment.



Note

When logging on to `servera` or `serverb`, you might see a message about activating `cockpit`. You can ignore the message.

```
[student@workstation ~]$ ssh student@serverb
Warning: Permanently added 'serverb,172.25.250.11' (ECDSA) to the list of
known hosts.
Activate the web console with: systemctl enable --now cockpit.socket

[student@serverb ~]$
```

Controlling Your Systems

You are assigned remote computers in a Red Hat Online Learning (ROLE) classroom. Self-paced courses are accessed through a web application that is hosted at `rol.redhat.com` [<http://rol.redhat.com>]. Log in to this site with your Red Hat Customer Portal user credentials.

Controlling the Virtual Machines

The virtual machines in your classroom environment are controlled through web page interface controls. The state of each classroom virtual machine is displayed on the **Lab Environment** tab.

Table of ContentsCourseLab Environment★i

▶ Lab Controls

Click **CREATE** to build all of the virtual machines needed for the classroom lab environment. This may take several minutes to complete. Once created the environment can then be stopped and restarted to pause your experience.

If you **DELETE** your lab, you will remove all of the virtual machines in your classroom and lose all of your progress.

DELETESTOPi

bastion	active	ACTION -	OPEN CONSOLE
classroom	active	ACTION -	OPEN CONSOLE
servera	building	ACTION -	OPEN CONSOLE
serverb	building	ACTION -	OPEN CONSOLE
workstation	active	ACTION -	OPEN CONSOLE

Figure 0.2: An example course Lab Environment management page

Machine States

Virtual Machine State	Description
building	The virtual machine is being created.
active	The virtual machine is running and available. If it just started, it might still be starting services.
stopped	The virtual machine is completely shut down. On starting, the virtual machine boots into the same state that it was in before shutdown. The disk state is preserved.

Classroom Actions

Button or Action	Description
CREATE	Create the ROLE classroom. Creates and starts all the needed virtual machines for this classroom. Creation can take several minutes to complete.
CREATING	The ROLE classroom virtual machines are being created. Creates and starts all the needed virtual machines for this classroom. Creation can take several minutes to complete.
DELETE	Delete the ROLE classroom. Destroys all virtual machines in the classroom. All saved work on those systems' disks is lost.
START	Start all virtual machines in the classroom.
STARTING	All virtual machines in the classroom are starting.
STOP	Stop all virtual machines in the classroom.

Machine Actions

Button or Action	Description
OPEN CONSOLE	Connect to the system console of the virtual machine in a new browser tab. You can log in directly to the virtual machine and run commands, when required. Normally, log in to the workstation virtual machine only, and from there, use ssh to connect to the other virtual machines.
ACTION > Start	Start (power on) the virtual machine.
ACTION > Shutdown	Gracefully shut down the virtual machine, preserving disk contents.
ACTION > Power Off	Forcefully shut down the virtual machine, while still preserving disk contents. This action is equivalent to removing the power from a physical machine.
ACTION > Reset	Forcefully shut down the virtual machine and reset associated storage to its initial state. All saved work on that system's disks is lost.

At the start of an exercise, if instructed to reset a single virtual machine node, then click **ACTION** > **Reset** for only that specific virtual machine.

At the start of an exercise, if instructed to reset all virtual machines, then click **ACTION** > **Reset** on every virtual machine in the list.

If you want to return the classroom environment to its original state at the start of the course, then click **DELETE** to remove the entire classroom environment. After the lab is deleted, then click **CREATE** to provision a new set of classroom systems.

**Warning**

The **DELETE** operation cannot be undone. All completed work in the classroom environment is lost.

The Auto-stop and Auto-destroy Timers

The Red Hat Online Learning enrollment entitles you to a set allotment of computer time. To help to conserve your allotted time, the ROLE classroom uses timers, which shut down or delete the classroom environment when the appropriate timer expires.

To adjust the timers, locate the two + buttons at the bottom of the course management page. Click the auto-stop + button to add another hour to the auto-stop timer. Click the auto-destroy + button to add another day to the auto-destroy timer. Auto-stop has a maximum of 11 hours, and auto-destroy has a maximum of 14 days. Be careful to keep the timers set while you are working, so that your environment is not unexpectedly shut down. Be careful not to set the timers unnecessarily high, which could waste your subscription time allotment.

Performing Lab Exercises

You might see the following lab activity types in this course:

- A *guided exercise* is a hands-on practice exercise that follows a presentation section. It walks you through a procedure to perform, step by step.
- A *quiz* is typically used when checking knowledge-based learning, or when a hands-on activity is impractical for some other reason.
- An *end-of-chapter lab* is a gradable hands-on activity to help you to check your learning. You work through a set of high-level steps, based on the guided exercises in that chapter, but the steps do not walk you through every command. A solution is provided with a step-by-step walk-through.
- A *comprehensive review lab* is used at the end of the course. It is also a gradable hands-on activity, and might cover content from the entire course. You work through a specification of what to accomplish in the activity, without receiving the specific steps to do so. Again, a solution is provided with a step-by-step walk-through that meets the specification.

To prepare your lab environment at the start of each hands-on activity, run the `lab start` command with a specified activity name from the activity's instructions. Likewise, at the end of each hands-on activity, run the `lab finish` command with that same activity name to clean up after the activity. Each hands-on activity has a unique name within a course.

The syntax for running an exercise script is as follows:

```
[student@workstation ~]$ lab action exercise
```

The *action* is a choice of `start`, `grade`, or `finish`. All exercises support `start` and `finish`. Only end-of-chapter labs and comprehensive review labs support `grade`.

start

The `start` action verifies the required resources to begin an exercise. It might include configuring settings, creating resources, checking prerequisite services, and verifying necessary outcomes from previous exercises. You can take an exercise at any time, even without taking preceding exercises.

grade

For gradable activities, the `grade` action directs the `lab` command to evaluate your work, and shows a list of grading criteria with a `PASS` or `FAIL` status for each. To achieve a `PASS` status for all criteria, fix the failures and rerun the `grade` action.

finish

The `finish` action cleans up resources that were configured during the exercise. You can take an exercise as many times as you want.

The `lab` command supports tab completion. For example, to list all exercises that you can start, enter `lab start` and then press the Tab key twice.

Chapter 1

Improve Command-line Productivity

Goal

Run commands more efficiently by using advanced features of the Bash shell, shell scripts, and various Red Hat Enterprise Linux utilities.

Objectives

- Run commands more efficiently by using advanced features of the Bash shell, shell scripts, and various Red Hat Enterprise Linux utilities.
- Run repetitive tasks with `for` loops, evaluate exit codes from commands and scripts, run tests with operators, and create conditional structures with `if` statements.
- Create regular expressions to match data, apply regular expressions to text files with the `grep` command, and use `grep` to search files and data from piped commands.

Sections

- Write Simple Bash Scripts (and Guided Exercise)
- Loops and Conditional Constructs in Scripts (and Guided Exercise)
- Match Text in Command Output with Regular Expressions (and Guided Exercise)

Lab

- Improve Command-line Productivity

Write Simple Bash Scripts

Objectives

- Run commands more efficiently by using advanced features of the Bash shell, shell scripts, and various Red Hat Enterprise Linux utilities.

Create and Execute Bash Shell Scripts

You can accomplish many system administration tasks by using command-line tools. More complex tasks often require chaining together multiple commands that pass results between them. By using the Bash shell environment and scripting features, you can combine Linux commands into *shell scripts* to solve repetitive real-world problems.

A Bash shell script is an executable file that contains a list of commands, and possibly with programming logic to control decision-making in the overall task. When well-written, a shell script is a powerful command-line tool on its own, and you can use it with other scripts.

Shell scripting proficiency is essential for system administrators in any operational environment. You can use shell scripts to improve the efficiency and accuracy of routine task completion.

Although you can use any text editor, advanced editors such as `vim` or `emacs` understand Bash shell syntax and can provide color-coded highlighting. This highlighting helps to identify common scripting errors such as improper syntax, unmatched quotes, parentheses, brackets, and braces, and other structural mistakes.

Specify the Command Interpreter

The first line of a script begins with the `#!` notation, which is commonly referred to as *she-bang* or *hash-bang*, from the names of those two characters, *sharp* or *hash* and *bang*. This notation is an interpreter directive to indicate the command interpreter and command options to process the remaining lines in the file. For Bash syntax script files, the first line is the following directive:

```
#!/usr/bin/bash
```

Execute a Bash Shell Script

A shell script file must have execute permissions to run it as an ordinary command. Use the `chmod` command to modify the file permissions. Use the `chown` command, if needed, to grant execute permission only for specific users or groups.

If the script is stored in a directory that is listed in the shell's `PATH` environmental variable, then you can run the shell script by using only its file name, similar to running compiled commands. Because `PATH` parsing runs the first matching file name that is found, always avoid using existing command names to name your script files. If a script is not in a `PATH` directory, then run the script by using its absolute path name, which you can determine by querying the file with the `which` command. Alternatively, run a script in your current working directory by using the `.` directory prefix, such as `./scriptname`.

```
[user@host ~]$ which hello
~/bin/hello
[user@host ~]$ echo $PATH
/home/user/.local/bin:/home/user/bin:/sbin:/bin:/usr/sbin:/usr/bin:/usr/local/
sbin:/usr/local/bin
```

Quote Special Characters

Some characters and words have a special meaning to the Bash shell. To use these characters for their literal values, rather than for their special meanings, you *escape* them in the script. Use the backslash character (`\`), single quotes (`' '`), or double quotes (`" "`) to remove (or *escape*) the special meaning of these characters.

The backslash character removes the special meaning of the single character that immediately follows the backslash. For example, to use the `echo` command to display the `# not a comment` literal string, the `#` hash character must not be interpreted as a comment.

The following example shows the backslash character (`\`) modifying the hash character so it is not interpreted as a comment:

```
[user@host ~]$ echo # not a comment

[user@host ~]$ echo \# not a comment
# not a comment
```

To escape more than one character in a text string, either use the backslash character multiple times, or enclose the whole string in single quotes (`' '`) to interpret literally. Single quotes preserve the literal meaning of all characters that they enclose. Observe the backslash character and single quotes in these examples:

```
[user@host ~]$ echo # not a comment #

[user@host ~]$ echo \# not a comment #
# not a comment
[user@host ~]$ echo \# not a comment \#
# not a comment #
[user@host ~]$ echo '# not a comment #'
# not a comment #
```

Use double quotation marks to suppress *globbing* (file name pattern matching) and shell expansion, but still allow command and variable substitution. Variable substitution is conceptually the same as command substitution, but might use optional brace syntax. Observe the following examples of various quotation mark forms.

Use single quotation marks to interpret all enclosed text literally. Besides suppressing globbing and shell expansion, single quotation marks also direct the shell to suppress command and variable substitution. The question mark (`?`) is included inside the quotations, because it is a *metacharacter* that also needs escaping from expansion.

```
[user@host ~]$ var=$(hostname -s); echo $var
host
[user@host ~]$ echo "***** hostname is ${var} *****"
***** hostname is host *****
```

```
[user@host ~]$ echo Your username variable is \${USER}.
Your username variable is $USER.
[user@host ~]$ echo "Will variable $var evaluate to ${hostname -s}?"
Will variable host evaluate to host?
[user@host ~]$ echo 'Will variable $var evaluate to ${hostname -s}?'
Will variable $var evaluate to ${hostname -s}?
[user@host ~]$ echo "\"Hello, world\""
"Hello, world"
[user@host ~]$ echo '"Hello, world"'
"Hello, world"
```

Provide Output from a Shell Script

The `echo` command displays arbitrary text by passing the text as an argument to the command. By default, the text is sent to *standard output* (*STDOUT*). You can send text elsewhere by using output redirection. In the following simple Bash script, the `echo` command displays the "Hello, world" message to *STDOUT*, which defaults to the screen device.

```
[user@host ~]$ cat ~/bin/hello
#!/usr/bin/bash

echo "Hello, world"

[user@host ~]$ hello
Hello, world
```



Note

This user can run `hello` at the prompt because the `~/bin` (`/home/user/bin`) directory is in the user's `PATH` variable and the `hello` script has executable permission. The `PATH` parser finds the script first, if no other executable file called `hello` is found in any earlier `PATH` directory. Your home directory's `bin` subdirectory is intended to store your personal scripts.

The `echo` command is widely used in shell scripts to display informational or error messages. Messages helpfully indicate a script's progress, and can be directed to standard output or standard error, or be redirected to a log file for archiving. When you display error messages, good programming practice is to redirect error messages to *STDERR* to separate them from normal program output.

```
[user@host ~]$ cat ~/bin/hello
#!/usr/bin/bash

echo "Hello, world"
echo "ERROR: Houston, we have a problem." >&2

[user@host ~]$ hello 2> hello.log
Hello, world
[user@host ~]$ cat hello.log
ERROR: Houston, we have a problem.
```

The `echo` command is also helpful to debug a problematic shell script. Adding `echo` statements in a script, to display variable values and other runtime information, can help to clarify how a script is behaving.

**References**

`bash(1)`, `echo(1)`, and `echo(1p)` man pages

► Guided Exercise

Write Simple Bash Scripts

Write a simple Bash script with a sequence of commands and run it from the command line.

Outcomes

- Write and execute a simple Bash script.
- Redirect the output of a simple Bash script to a file.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start console-write
```

Instructions

- 1. Log in to the servera machine as the student user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
```

- 2. Create and execute a simple Bash script.

- 2.1. Use the `vim` command to create the `firstscript.sh` file under your home directory.

```
[student@servera ~]$ vim firstscript.sh
```

- 2.2. Insert the following text, and save the file. The number of hash signs (`#`) is arbitrary.

```
#!/usr/bin/bash
echo "This is my first bash script" > ~/output.txt
echo "" >> ~/output.txt
echo "#####" >> ~/output.txt
```

- 2.3. Use the `bash` command to execute the script.

```
[student@servera ~]$ bash firstscript.sh
```

- 2.4. Review the output file that the script generated.

```
[student@servera ~]$ cat output.txt
This is my first bash script
```

```
#####
```

- 3. Add more commands to the `firstscript.sh` script, execute it, and review the output.

3.1. Use the Vim text editor to edit the `firstscript.sh` script.

```
[student@servera ~]$ vim firstscript.sh
```

The following output shows the expected content of the `firstscript.sh` file:

```
#!/usr/bin/bash
#
echo "This is my first bash script" > ~/output.txt
echo "" >> ~/output.txt
echo "#####" >> ~/output.txt
echo "LIST BLOCK DEVICES" >> ~/output.txt
echo "" >> ~/output.txt
lsblk >> ~/output.txt
echo "" >> ~/output.txt
echo "#####" >> ~/output.txt
echo "FILESYSTEM FREE SPACE STATUS" >> ~/output.txt
echo "" >> ~/output.txt
df -h >> ~/output.txt
echo "#####" >> ~/output.txt
```

3.2. Make the `firstscript.sh` file executable by using the `chmod` command.

```
[student@servera ~]$ chmod a+x firstscript.sh
```

3.3. Execute the `firstscript.sh` script.

```
[student@servera ~]$ ./firstscript.sh
```

3.4. Review the output file that the script generated.

```
[student@servera ~]$ cat output.txt
This is my first bash script

#####
LIST BLOCK DEVICES

NAME      MAJ:MIN RM  SIZE RO TYPE MOUNTPOINTS
sr0        11:0    1   558K  0 rom
vda       252:0    0    10G  0 disk
└─vda1    252:1    0     1M  0 part
└─vda2    252:2    0   200M  0 part /boot/efi
└─vda3    252:3    0   500M  0 part /boot
└─vda4    252:4    0   9.3G  0 part /
```

```
vdb    252:16    0    5G    0 disk
vdc    252:32    0    5G    0 disk
vdd    252:48    0    5G    0 disk

#####
FILESYSTEM FREE SPACE STATUS

Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        844M    0  844M   0% /dev
tmpfs           888M    0  888M   0% /dev/shm
tmpfs           355M  9.4M  346M   3% /run
/dev/vda4       9.4G  1.7G  7.7G  18% /
/dev/vda3       495M  161M  335M  33% /boot
/dev/vda2       200M   7.6M  193M   4% /boot/efi
tmpfs           178M    0  178M   0% /run/user/1000
#####
```

- 4. Remove the exercise files and return to the workstation machine.

4.1. Delete the `firstscript.sh` and `output.txt` files.

```
[student@servera ~]$ rm firstscript.sh output.txt
```

4.2. Return to the workstation machine as the student user.

```
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the workstation machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish console-write
```

Loops and Conditional Constructs in Scripts

Objectives

- Run repetitive tasks with `for` loops, evaluate exit codes from commands and scripts, run tests with operators, and create conditional structures with `if` statements.

Use Loops to Iterate Commands

System administrators often encounter repetitive tasks in their daily activities. A repetitive task example is running a command multiple times on a target, such as checking a process every minute for 10 minutes to know whether it has completed. Another example is running a command once each for multiple targets, such as backing up many databases on a system. The *for loop* is a Bash looping construct to use for task iterations.

Process Items from the Command Line

In Bash, the `for` loop construct uses the following syntax:

```
for VARIABLE in LIST; do
  COMMAND VARIABLE
done
```

The loop processes the strings that you provide in *LIST* and exits after processing the last string in the list. The `for` loop temporarily stores each list string as the value of *VARIABLE*, and then executes the block of commands that use the variable. The variable name is arbitrary. Typically, you reference the variable value with commands in the command block.

Provide the list of strings for the `for` loop from a list that the user enters directly, or that is generated from shell expansion, such as variable, brace, or file name expansion, or command substitution.

These examples demonstrate different ways to provide strings to `for` loops:

```
[user@host ~]$ for HOST in host1 host2 host3; do echo $HOST; done
host1
host2
host3
[user@host ~]$ for HOST in host{1,2,3}; do echo $HOST; done
host1
host2
host3
[user@host ~]$ for HOST in host{1..3}; do echo $HOST; done
host1
host2
host3
[user@host ~]$ for FILE in file{a..c}; do ls $FILE; done
filea
fileb
filec
```

```
[user@host ~]$ for PACKAGE in $(rpm -qa | grep kernel); \
do echo "$PACKAGE was installed on \
$(date -d @$(rpm -q --qf "%{INSTALLTIME}\n" $PACKAGE)); done
kernel-tools-libs-5.14.0-70.2.1.el9_0.x86_64 was installed on Thu Mar 24 10:52:40
PM EDT 2022
kernel-tools-5.14.0-70.2.1.el9_0.x86_64 was installed on Thu Mar 24 10:52:40 PM
EDT 2022
kernel-core-5.14.0-70.2.1.el9_0.x86_64 was installed on Thu Mar 24 10:52:46 PM EDT
2022
kernel-modules-5.14.0-70.2.1.el9_0.x86_64 was installed on Thu Mar 24 10:52:47 PM
EDT 2022
kernel-5.14.0-70.2.1.el9_0.x86_64 was installed on Thu Mar 24 10:53:04 PM EDT 2022
[user@host ~]$ for EVEN in $(seq 2 2 10); do echo "$EVEN"; done
2
4
6
8
10
```

Bash Script Exit Codes

After a script interprets and processes all of its content, the script process exits and passes back control to the parent process that called it. However, a script can be exited before it finishes, such as when the script encounters an error condition. Use the `exit` command to immediately leave the script, and skip processing the remainder of the script.

Use the `exit` command with an optional integer argument between 0 and 255, which represents an *exit code*. An exit code is returned to a parent process to indicate the status at exit. An exit code value of 0 represents a successful script completion with no errors. All other nonzero values indicate an error exit code. The script programmer defines these codes. Use unique values to represent the different error conditions that are encountered. Retrieve the exit code of the last completed command from the built-in `$?` variable, as in the following examples:

```
[user@host bin]$ cat hello
#!/usr/bin/bash
echo "Hello, world"
exit 0
[user@host bin]$ ./hello
Hello, world
[user@host bin]$ echo $?
0
```

When a script's `exit` command is used without an exit code argument, the script returns the exit code of the last command that was run within the script.

Test Logic for Strings and Directories, and to Compare Values

To ensure that unexpected conditions do not disrupt scripts, it is recommended to verify command input such as command-line arguments, user input, command substitutions, variable expansions, and file name expansions. You can verify integrity in your scripts by using the Bash `test` command.

All commands produce an exit code on completion.

To see the exit status, view the `$?` variable immediately after executing the `test` command. An exit status of 0 indicates a successful exit with nothing to report. Nonzero values indicate some condition or failure. Use various operators to test whether a number is greater than (`gt`), greater than or equal to (`ge`), less than (`lt`), less than or equal to (`le`), or equal (`eq`) to another number.

Use operators to test whether a string of text is the same (`=` or `==`) or not the same (`!=`) as another string of text, or whether the string has zero length (`z`) or has a non-zero length (`n`). You can also test whether a regular file (`-f`) or directory (`-d`) exists, and has some special attributes, such as if the file is a symbolic link (`-L`), or if the user has read permissions (`-r`).

**Note**

Shell scripting uses many other operator types. The `test(1)` man page lists the conditional expression operators with descriptions. The `bash(1)` man page also explains operator use and evaluation, but can be complex to read. Red Hat recommends learning shell scripting through quality books and courses that are dedicated to shell programming.

The following examples demonstrate the `test` command with Bash numeric comparison operators:

```
[user@host ~]$ test 1 -gt 0 ; echo $?
0
[user@host ~]$ test 0 -gt 1 ; echo $?
1
```

Test by using the Bash `test` command syntax, `[ <TESTEXPRESSION> ]` or the newer extended `test` command syntax, `[[ <TESTEXPRESSION> ]]`, which provides features such as file name globbing and regex pattern matching. In most cases, use the `[[ <TESTEXPRESSION> ]]` syntax.

The following examples demonstrate the Bash `test` command syntax and numeric comparison operators:

```
[user@host ~]$ [[ 1 -eq 1 ]]; echo $?
0
[user@host ~]$ [[ 1 -ne 1 ]]; echo $?
1
[user@host ~]$ [[ 8 -gt 2 ]]; echo $?
0
[user@host ~]$ [[ 2 -ge 2 ]]; echo $?
0
[user@host ~]$ [[ 2 -lt 2 ]]; echo $?
1
[user@host ~]$ [[ 1 -lt 2 ]]; echo $?
0
```

The following examples demonstrate the Bash string comparison operators:


```
[user@host ~]$ [[ abc = abc ]]; echo $?
0
[user@host ~]$ [[ abc == def ]]; echo $?
1
[user@host ~]$ [[ abc != def ]]; echo $?
0
```

The following examples demonstrate Bash string unary (one argument) operators:

```
[user@host ~]$ STRING=''; [[ -z "$STRING" ]]; echo $?
0
[user@host ~]$ STRING='abc'; [[ -n "$STRING" ]]; echo $?
0
```



Note

The space characters inside the brackets are mandatory, because they separate the words and elements within the test expression. The shell's command parsing routine divides the command elements into words and operators by recognizing spaces and other metacharacters, according to built-in parsing rules. For full treatment of this advanced concept, see the `getopt(3)` man page. The left square bracket character (`[`) is itself a built-in alias for the `test` command. Shell words, whether they are commands, subcommands, options, arguments, or other token elements, are always delimited by spaces.

Conditional Structures

Simple shell scripts represent a collection of commands that are executed from beginning to end. Programmers incorporate decision-making into shell scripts by using conditional structures. A script can execute specific routines when stated conditions are met.

Use the If/Then Construct

The simplest conditional structure is the *if/then* construct, with the following syntax:

```
if <CONDITION>; then
    <STATEMENT>
    ...
    <STATEMENT>
fi
```

With this construct, if the script meets the given condition, then it executes the code in the statement block. It does not act if the given condition is not met. Common test conditions in the `if/then` statements include the previously discussed numeric, string, and file tests. The `fi` statement at the end closes the `if/then` construct. The following code section demonstrates an `if/then` construct to start the `psacct` service if it is not active:

```
[user@host ~]$ systemctl is-active psacct > /dev/null 2>&1
[user@host ~]$ if [[ $? -ne 0 ]]; then sudo systemctl start psacct; fi
```

Use the If/Then/Else Construct

You can further expand the `if/then` construct to take different sets of actions depending on whether a condition is met. Use the `if/then/else` construct to accomplish this behavior, as in this example:

```
if <CONDITION>; then
    <STATEMENT>
    ...
    <STATEMENT>
else
    <STATEMENT>
    ...
    <STATEMENT>
fi
```

The following code section demonstrates an `if/then/else` statement to start the `psacct` service if it is not active, and to stop it if it is active:

```
[user@host ~]$ systemctl is-active psacct > /dev/null 2>&1
[user@host ~]$ if [[ $? -ne 0 ]]; then \
sudo systemctl start psacct; \
else \
sudo systemctl stop psacct; \
fi
```

Use the If/Then/Elif/Then/Else Construct

Expand an `if/then/else` construct to test more than one condition and to execute a different set of actions when it meets a specific condition. The next example shows the construct for an added condition:

```
if <CONDITION>; then
    <STATEMENT>
    ...
    <STATEMENT>
elif <CONDITION>; then
    <STATEMENT>
    ...
    <STATEMENT>
else
    <STATEMENT>
    ...
    <STATEMENT>
fi
```

In this conditional structure, Bash tests the conditions as they are ordered in the script. When a condition is true, Bash executes the actions that are associated with the condition and then skips the remainder of the conditional structure. If none of the conditions are true, then Bash executes the actions in the `else` clause.

The following example demonstrates an `if/then/elif/then/else` statement to run the `mysql` client if the `mariadb` service is active, or to run the `psql` client if the `postgresql` service

is active, or to run the `sqlite3` client if both the `mariadb` and the `postgresql` service are inactive:

```
[user@host ~]$ systemctl is-active mariadb > /dev/null 2>&1
[user@host ~]$ MARIADB_ACTIVE=$?
[user@host ~]$ sudo systemctl is-active postgresql > /dev/null 2>&1
[user@host ~]$ POSTGRESQL_ACTIVE=$?
[user@host ~]$ if [[ "$MARIADB_ACTIVE" -eq 0 ]]; then \
mysql; \
elif [[ "$POSTGRESQL_ACTIVE" -eq 0 ]]; then \
psql; \
else \
sqlite3; \
fi
```



References

bash(1) man page

► Guided Exercise

Loops and Conditional Constructs in Scripts

Use loops to efficiently print the hostname from multiple servers.

Outcomes

- Create a `for` loop to iterate through a list of items from the command line and in a shell script.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start console-commands
```

Instructions

- 1. Use the `ssh` and `hostname` commands to print the hostname of the `servera` and `serverb` machines to standard output.

```
[student@workstation ~]$ ssh student@servera hostname
servera.lab.example.com
[student@workstation ~]$ ssh student@serverb hostname
serverb.lab.example.com
```

- 2. Create a `for` loop to execute the `hostname` command on the `servera` and `serverb` machines.

```
[student@workstation ~]$ for HOST in servera serverb
do
ssh student@${HOST} hostname
done
servera.lab.example.com
serverb.lab.example.com
```

- 3. Create a shell script in the `/home/student/bin` directory to execute the same `for` loop. Ensure that the script is included in the `PATH` environment variable.
 - 3.1. Create the `/home/student/bin` directory to store the shell script, if the directory does not exist.

```
[student@workstation ~]$ mkdir ~/bin
```

- 3.2. Verify that the `bin` subdirectory of your home directory is in your `PATH` environment variable.

```
[student@workstation ~]$ echo $PATH
/home/student/.local/bin:/home/student/bin:/sbin:/bin:/usr/sbin:/usr/bin:/usr/local/sbin:/usr/local/bin:/home/student/.venv/labs/bin
```

- 3.3. Create a shell script called `printhostname.sh` in the `/home/student/bin` directory to perform the `for` loop, and add the following content in the file.

```
[student@workstation ~]$ vim ~/bin/printhostname.sh
#!/usr/bin/bash
#Execute for loop to print server hostname.
for HOST in servera serverb
do
    ssh student@${HOST} hostname
done
exit 0
```

- 3.4. Give the created script executable permission.

```
[student@workstation ~]$ chmod +x ~/bin/printhostname.sh
```

- 3.5. Run the script from your home directory.

```
[student@workstation ~]$ printhostname.sh
servera.lab.example.com
serverb.lab.example.com
```

- 3.6. Verify that the exit code of your script is 0.

```
[student@workstation ~]$ echo $?
0
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish console-commands
```

Match Text in Command Output with Regular Expressions

Objectives

- Create regular expressions to match data, apply regular expressions to text files with the `grep` command, and use `grep` to search files and data from piped commands.

Write Regular Expressions

Regular expressions provide a pattern matching mechanism to find specific content. The `vim`, `grep`, and `less` commands can use regular expressions. Programming languages such as Perl, Python, and C also support regular expressions, but might differ slightly in syntax.

Regular expressions are a unique language, with their own syntax and rules. This section introduces regular expression syntax as implemented in `bash`, with examples.

Describe a Simple Regular Expression

The simplest regular expression is an exact match of the string to search. An exact match is when the characters in the regular expression match the type and order of the string.

Imagine that a user is looking through the following file for all occurrences of the pattern `cat`:

```
cat
dog
concatenate
dogma
category
educated
boondoggle
vindication
chilidog
```

The `cat` string is an exact match of the `c` character, followed by the `a` and `t` characters with no other characters between. Searching the file with the `cat` string as the regular expression returns the following matches:

```
cat
concatenate
category
educated
vindication
```

Match the Start and End of a Line

The regular expression would match the search string anywhere on the line on which it occurred: the beginning, middle, or end of the word or line. Use a *line anchor* metacharacter to control where on a line to look for a match.

To match only at the beginning of a line, use the caret character (^). To match only at the end of a line, use the dollar sign (\$).

With the same file as for the previous example, the `^cat` regular expression would match two lines.

```
cat
category
```

The `cat$` regular expression would find only one match, where the `cat` characters are at the end of a line.

```
cat
```

Locate lines in the file that end with `dog`, by using an end-of-line anchor to create the `dog$` regular expression, which matches two lines:

```
dog
chilidog
```

To locate a line that contains only the search expression exactly, use both the beginning and end-of-line anchors. For example, to locate the word `cat` when it is both at the beginning and the end of a line simultaneously, use `^cat$`.

```
cat
```

Basic and Extended Regular Expression

The two types of regular expressions are basic regular expressions and extended regular expressions.

One difference between basic and extended regular expressions is in the behavior of the `|`, `+`, `?`, `(`, `)`, `{`, and `}` special characters. In basic regular expression syntax, these characters have a special meaning only if they are prefixed with a backslash `\` character. In extended regular expression syntax, these characters are special unless they are prefixed with a backslash `\` character. Other minor differences apply to how the `^`, `$`, and `*` characters are handled.

The `grep`, `sed`, and `vim` commands use basic regular expressions. The `grep` command `-E` option, the `sed` command `-E` option, and the `less` command use extended regular expressions.

Wildcard and Multiplier Usage in Regular Expressions

Regular expressions use a dot character (`.`) as a wildcard to match any single character on a single line. The `c.t` regular expression searches for a string that contains a `c`, followed by any single character, followed by a `t`. Example matches might include `cat`, `concatenate`, `vindication`, `cut`, and `c$t`.

With an unrestricted wildcard, you cannot predict the character that matches the wildcard. To match specific characters, replace the unrestricted wildcard with appropriate characters.

The use of bracket characters, such as in the `c[aou]t` regular expression, matches patterns that start with a `c`, followed by an `a`, `o`, or `u`, followed by a `t`. Possible matching expressions can have the `cat`, `cot`, and `cut` strings.

Multipliers are an often used mechanism with wildcards. Multipliers apply to the previous character or wildcard in the regular expression. An often used multiplier is the asterisk (*) character. When used in a regular expression, the asterisk multiplier matches zero or more occurrences of the multiplied expression. You can use the asterisk with expressions, in addition to characters.

For example, the `c[aou]*t` regular expression might match `coat` or `coot`. A regular expression of `c.*t` matches `cat`, `coat`, `culvert`, and even `ct` (matching zero characters between the `c` and the `t`). Any string that starts with a `c`, is followed by zero or more characters, and ends with a `t` must be a match.

Another type of multiplier indicates a more precise number of characters in the pattern. An example of an explicit multiplier is the `'c.\{2\}t'` regular expression, which matches any word that begins with a `c`, followed by exactly any two characters, and ends with a `t`. The `'c.\{2\}t'` expression would match two words in the following example:

```
cat
coat
convert
cart
covert
cypher
```



Note

This course introduced two metacharacter text parsing mechanisms: shell *pattern matching* (also known as file globbing or file name expansion), and *regular expressions*. Both mechanisms use similar metacharacters, such as the asterisk character (*), but have differences in metacharacter interpretation and rules.

Pattern matching is a shell technique to specify multiple file names on the command line. Regular expressions represent any form or pattern in text strings, no matter how complex. Regular expressions are internally supported by many text processing commands, such as `grep`, `sed`, `awk`, `python`, and `perl`, and in many applications.

Basic and Extended Regular Expression Syntax

Basic syntax	Extended syntax	Description
.		The period (.) matches any single character.
?		The preceding item is optional and is matched at most once.
*		The preceding item is matched zero or more times.
+		The preceding item is matched one or more times.
\{n\}	{n}	The preceding item is matched exactly n times.
\{n,\}	{n,}	The preceding item is matched n or more times.
\{,m\}	{,m}	The preceding item is matched at most m times.

Basic syntax	Extended syntax	Description
<code>\{n,m\}</code>	<code>{n,m}</code>	The preceding item is matched at least <i>n</i> times, but not more than <i>m</i> times.
<code>[:alnum:]</code>		Alphanumeric characters: <code>[:alpha:]</code> and <code>[:digit:]</code> ; in the 'C' locale and ASCII character encoding, this expression is the same as <code>[0-9A-Za-z]</code> .
<code>[:alpha:]</code>		Alphabetic characters: <code>[:lower:]</code> and <code>[:upper:]</code> ; in the 'C' locale and ASCII character encoding, this expression is the same as <code>[A-Za-z]</code> .
<code>[:blank:]</code>		Blank characters: space and tab.
<code>[:cntrl:]</code>		Control characters. In ASCII, these characters have octal codes 000 through 037, and 177 (DEL).
<code>[:digit:]</code>		Digits: 0 1 2 3 4 5 6 7 8 9.
<code>[:graph:]</code>		Graphical characters: <code>[:alnum:]</code> and <code>[:punct:]</code> .
<code>[:lower:]</code>		Lowercase letters; in the 'C' locale and ASCII character encoding: a b c d e f g h i j k l m n o p q r s t u v w x y z.
<code>[:print:]</code>		Printable characters: <code>[:alnum:]</code> , <code>[:punct:]</code> , and space.
<code>[:punct:]</code>		Punctuation characters; in the 'C' locale and ASCII character encoding: ! " # \$ % & ' () * + , - . / : ; < = > ? @ [\] ^ _ ` { } ~.
<code>[:space:]</code>		Space characters: in the 'C' locale, it is tab, newline, vertical tab, form feed, carriage return, and space.
<code>[:upper:]</code>		Uppercase letters: in the 'C' locale and ASCII character encoding, it is: A B C D E F G H I J K L M N O P Q R S T U V W X Y Z.
<code>[:xdigit:]</code>		Hexadecimal digits: 0 1 2 3 4 5 6 7 8 9 A B C D E F a b c d e f.
<code>\b</code>		Match the empty string at the edge of a word.
<code>\B</code>		Match the empty string provided that it is not at the edge of a word.
<code>\<</code>		Match the empty string at the beginning of a word.
<code>\></code>		Match the empty string at the end of a word.
<code>\w</code>		Match word constituent. Synonym for <code>[_[:alnum:]]</code> .
<code>\W</code>		Match non-word constituent. Synonym for <code>[^_[:alnum:]]</code> .
<code>\s</code>		Match white space. Synonym for <code>[:space:]</code> .

Basic syntax	Extended syntax	Description
\S		Match non-white space. Synonym for <code>[^[:space:]]</code> .

Match Regular Expressions from the Command Line

The `grep` command uses regular expressions to isolate matching data. You can use the `grep` command to match data in a single file or in multiple files. When you use `grep` to match data in multiple files, it prints the file name followed by a colon character and then the lines that match the regular expression.

Isolating Data with the `grep` Command

The `grep` command specifies a regular expression and a file to parse for matches.

```
[user@host ~]$ grep '^computer' /usr/share/dict/words
computer
computerese
computerise
computerite
computerizable
computerization
computerize
computerized
computerizes
computerizing
computerlike
computernik
computers
```



Note

It is recommended practice to use single quotation marks to encapsulate the regular expression to protect any shell metacharacters (such as the `$`, `*`, and `{}` characters). Encapsulating the regular expression ensures that the command and not the shell interprets the characters.

The `grep` command can process output from other commands by using a pipe operator character (`|`). The following example shows the `grep` command parsing lines from the output of another command.

```
[root@host ~]# ps aux | grep chrony
chrony      662  0.0  0.1 29440 2468 ?        S    10:56   0:00 /usr/sbin/chronyd
```

The `grep` Command Options

The `grep` command has many options for controlling how it parses lines.

Table of Common grep Options

Option	Function
-i	Use the provided regular expression and do not enforce case sensitivity (run case-insensitive).
-v	Display only lines that do <i>not</i> contain matches to the regular expression.
-r	Search for data that matches the regular expression recursively in a group of files or directories.
-A <i>NUMBER</i>	Display <i>NUMBER</i> of lines after the regular expression match.
-B <i>NUMBER</i>	Display <i>NUMBER</i> of lines before the regular expression match.
-e	If multiple -e options are used, then multiple regular expressions can be supplied and are used with a logical OR.
-E	Use extended regular expression syntax instead of basic regular expression syntax when parsing the provided regular expression.

View the man pages to find other options for the `grep` command.

Examples of the grep Command

The following examples use various configuration files and log files.

Regular expressions are case-sensitive by default. Use the `grep` command -i option to run a case-insensitive search. The following example shows an excerpt of the `/etc/httpd/conf/httpd.conf` configuration file.

```
[user@host ~]$ cat /etc/httpd/conf/httpd.conf
...output omitted...
ServerRoot "/etc/httpd"

#
# Listen: Allows you to bind Apache to specific IP addresses and/or
# ports, instead of the default. See also the <VirtualHost>
# directive.
#
# Change this to Listen on a specific IP address, but note that if
# httpd.service is enabled to run at boot time, the address may not be
# available when the service starts. See the httpd.service(8) man
# page for more information.
#
#Listen 12.34.56.78:80
Listen 80
...output omitted...
```

The following example searches for the `serverroot` regular expression in the `/etc/httpd/conf/httpd.conf` configuration file.

```
[user@host ~]$ grep -i serverroot /etc/httpd/conf/httpd.conf
# with "/", the value of ServerRoot is prepended -- so 'log/access_log'
# with ServerRoot set to '/www' will be interpreted by the
# ServerRoot: The top of the directory tree under which the server's
# ServerRoot at a non-local disk, be sure to specify a local disk on the
# same ServerRoot for multiple httpd daemons, you will need to change at
ServerRoot "/etc/httpd"
```

Use the `grep` command `-v` option to *reverse* search the regular expression. This option displays only the lines that do *not* match the regular expression.

In the following example, all lines, regardless of case, that *do not* contain the `server` regular expression are returned.

```
[user@host ~]$ grep -v -i server /etc/hosts
127.0.0.1 localhost.localdomain localhost
172.25.254.254 classroom.example.com classroom
172.25.254.254 content.example.com content
172.25.254.254 materials.example.com materials
### rht-vm-hosts file listing the entries to be appended to /etc/hosts

172.25.250.9 workstation.lab.example.com workstation
172.25.250.254 bastion.lab.example.com bastion
172.25.250.220 utility.lab.example.com utility
172.25.250.220 registry.lab.example.com registry
```

To view a file without the distraction of comment lines, use the `grep` command `-v` option. In the following example, the regular expression matches and excludes all the lines that begin with a hash character (`#`) or a semicolon (`;`) character in the `/etc/systemd/system/multi-user.target.wants/rsyslog.service` file. In that file, the hash character at the beginning of a line indicates a general comment, whereas the semicolon character refers to a commented variable value.

```
[user@host ~]$ grep -v '^[#;]' \
/etc/systemd/system/multi-user.target.wants/rsyslog.service
[Unit]
Description=System Logging Service
Documentation=man:rsyslogd(8)
Documentation=https://www.rsyslog.com/doc/

[Service]
Type=notify
EnvironmentFile=/etc/sysconfig/rsyslog
ExecStart=/usr/sbin/rsyslogd -n $SYSLOGD_OPTIONS
ExecReload=/usr/bin/kill -HUP $MAINPID
UMask=0066
StandardOutput=null
Restart=on-failure

LimitNOFILE=16384
```

```
[Install]
WantedBy=multi-user.target
```

The `grep` command `-e` option can search for more than one regular expression at a time. The following example, which uses a combination of the `less` and `grep` commands, locates all occurrences of `pam_unix`, `user root`, and `Accepted publickey` in the `/var/log/secure` log file.

```
[root@host ~]# cat /var/log/secure | grep -e 'pam_unix' \
-e 'user root' -e 'Accepted publickey' | less
Mar  4 03:31:41 localhost passwd[6639]: pam_unix(passwd:chauthtok): password
  changed for root
Mar  4 03:32:34 localhost sshd[15556]: Accepted publickey for devops from
  10.30.0.167 port 56472 ssh2: RSA SHA256:M8ikhcEDm2tQ95Z0o7ZvufqEixCFct
+wowZLNzNlBT0
Mar  4 03:32:34 localhost systemd[15560]: pam_unix(systemd-user:session): session
  opened for user devops(uid=1001) by (uid=0)
```

To search for text in a file that you opened with the `vim` or `less` commands, first enter the slash character (`/`) and then type the pattern to find. Press `Enter` to start the search. Press `N` to find the next match.

```
[root@host ~]# vim /var/log/boot.log
...output omitted...
^[0;32m OK ^[[0m Finished ^[[0;1;39mdracut pre-pivot and cleanup hook^[0m.^M
  Starting ^[[0;1;39mCleaning Up and Shutting Down Daemons^[0m...^M
^[0;32m OK ^[[0m Stopped target ^[[0;1;39mRemote Encrypted Volumes^[0m.^M
^[0;32m OK ^[[0m Stopped target ^[[0;1;39mTimer Units^[0m.^M
^[0;32m OK ^[[0m Closed ^[[0;1;39mD-Bus System Message Bus Socket^[0m.^M
/Daemons
```

```
[root@host ~]# less /var/log/messages
...output omitted...
Mar  4 03:31:19 localhost kernel: pci 0000:00:02.0: vgaarb: setting as boot VGA
  device
Mar  4 03:31:19 localhost kernel: pci 0000:00:02.0: vgaarb: VGA device added:
  decodes=io+mem,owns=io+mem,locks=None
Mar  4 03:31:19 localhost kernel: pci 0000:00:02.0: vgaarb: bridge control
  possible
Mar  4 03:31:19 localhost kernel: vgaarb: loaded
Mar  4 03:31:19 localhost kernel: SCSI subsystem initialized
Mar  4 03:31:19 localhost kernel: ACPI: bus type USB registered
Mar  4 03:31:19 localhost kernel: usbcore: registered new interface driver usbfs
Mar  4 03:31:19 localhost kernel: usbcore: registered new interface driver hub
Mar  4 03:31:19 localhost kernel: usbcore: registered new device driver usb
/device
```



References

`regex(7)` and `grep(1)` man pages

► Guided Exercise

Match Text in Command Output with Regular Expressions

Search for text in the system logs and the output of commands to find information more efficiently.

Outcomes

- Efficiently search for text in log files and configuration files.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command ensures that all required resources are available.

```
[student@workstation ~]$ lab start console-regex
```

Instructions

- 1. Log in to the servera machine as the student user and switch to the root user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

- 2. Use the `grep` command to find the GID and UID for the `postfix` and `postdrop` groups and users. To do so, use the `rpm -q --scripts` command, which queries the information for a specific package and shows the scripts that are used as part of the installation process.

```
[root@servera ~]# rpm -q --scripts postfix | grep -e 'user' -e 'group'
# Add user and groups if necessary
# generated from postfix.sysusers
getent group 'postfix' >/dev/null || groupadd -f -g '89' -r 'postfix' || :
    useradd -r -u '89' -g 'postfix' -d '/var/spool/postfix' -s '/sbin/nologin'
-c '' 'postfix' || :
    useradd -r -g 'postfix' -d '/var/spool/postfix' -s '/sbin/nologin' -c ''
'postfix' || :
getent group 'postdrop' >/dev/null || groupadd -f -g '90' -r 'postdrop' || :
getent group 'mail' >/dev/null || groupadd -f -g '12' -r 'mail' || :
getent group 'mail' >/dev/null || groupadd -r 'mail' || :
    useradd -r -g 'postfix' -d '/' -s '/usr/sbin/nologin' -c '' 'postfix' || :
```

```
if getent group 'mail' >/dev/null; then
    usermod -a -G 'mail' 'postfix' || :
    setgid_group=postdrop \
```

- ▶ 3. Modify the previous regular expression to display the first two messages in the `/var/log/maillog` file. In this search, you do not need to use the caret character (^), because you are not searching for the first character in a line.

```
[root@servera ~]# grep 'postfix' /var/log/maillog | head -n 2
Apr  1 15:27:16 servera postfix/postfix-script[3121]: starting the Postfix mail
system
Apr  1 15:27:16 servera postfix/master[3123]: daemon started -- version 3.5.9,
configuration /etc/postfix
```

- ▶ 4. Find the name of the queue directory for the Postfix server. Search the `/etc/postfix/main.cf` configuration file for all information about queues. Use the `grep` command `-i` option to ignore case distinctions.

```
[root@servera ~]# grep -i 'queue' /etc/postfix/main.cf
# testing.  When soft_bounce is enabled, mail will remain queued that
# The queue_directory specifies the location of the Postfix queue.
queue_directory = /var/spool/postfix
# QUEUE AND PROCESS OWNERSHIP
# The mail_owner parameter specifies the owner of the Postfix queue
# is the Sendmail-compatible mail queue listing command.
# setgid_group: The group for mail submission and queue management
```

- ▶ 5. Confirm that the postfix service writes messages to the `/var/log/messages` file. Use the `less` command and then the slash character (/) to search the file. Press `n` to move to the next entry that matches the search. Press `q` to quit the `less` command.

```
[root@servera ~]# less /var/log/messages
...output omitted...
Apr  1 15:27:15 servera systemd[1]: Starting Postfix Mail Transport Agent...
...output omitted...
Apr  1 15:27:16 servera systemd[1]: Started Postfix Mail Transport Agent.
...output omitted...
/Postfix
```

- ▶ 6. Use the `ps aux` command to confirm that the postfix server is currently running. Use the `grep` command to limit the output to the necessary lines.

```
[root@servera ~]# ps aux | grep postfix
root      3123  0.0  0.2 38172 4384 ?        Ss   15:27   0:00 /usr/
libexec/postfix/master -w
postfix   3124  0.0  0.4 45208 8236 ?        S    15:27   0:00 pickup -l -t
unix -u
postfix   3125  0.0  0.4 45252 8400 ?        S    15:27   0:00 qmgr -l -t unix
-u
root      3228  0.0  0.1 221668 2288 pts/0    S+   15:55   0:00 grep --
color=auto postfix
```

- 7. Confirm that the `qmgr`, `cleanup`, and `pickup` queues are correctly configured. Use the `grep` command `-e` option to match multiple entries in the same file. The `/etc/postfix/master.cf` file is the configuration file.

```
[root@servera ~]# grep -e qmgr -e pickup -e cleanup /etc/postfix/master.cf
pickup      unix  n      -      n      60      1      pickup
cleanup     unix  n      -      n      -       0      cleanup
qmgr        unix  n      -      n      300     1      qmgr
#qmgr       unix  n      -      n      300     1      oqmgr
```

- 8. Return to the `workstation` machine as the `student` user.

```
[root@servera ~]# exit
logout
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish console-regex
```

► Lab

Improve Command-line Productivity

Create a Bash script that can filter and get relevant information from different hosts.

Outcomes

- Create a Bash script and redirect its output to a file.
- Use loops to simplify your code.
- Filter the relevant content by using `grep` and regular expressions.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start console-review
```

Instructions

1. Create the executable `/home/student/bin/bash-lab` script file on the `workstation` machine. The initial content in the script must use the shebang interpreter directive.
2. Edit your newly created script file to store the following information from the `servera` and `serverb` machines on the `workstation` machine. The systems use SSH keys for authentication, and therefore you do not require a password. Store the output of the listed commands from the following table in the `/home/student/output-servera` and `/home/student/output-serverb` files respectively on the `workstation` machine. Use the hash sign (`#`) for differentiating the output of the successive commands in the output file.

Command or file	Content requested
hostname -f	Store the entire output.
echo "#####"	Append the hash signs to differentiate the following command.
lscpu	Get only the lines that start with the CPU string.
echo "#####"	Append the hash signs to differentiate the following command.
/etc/selinux/config	Ignore empty lines. Also, ignore lines that start with the # character.
echo "#####"	Append the hash signs to differentiate the following command.
/var/log/secure	Get all "Failed password" entries.
echo "#####"	Append the hash signs to differentiate the following command.

Save the required information to the `output-servera` and `output-serverb` files in the `/home/student` directory on `workstation`.



Note

You can use the `sudo` command without requiring a password on the `servera` and `serverb` hosts. Remember to use a loop to simplify your script. You can also use multiple `grep` commands that are concatenated with the use of the pipe character (`|`).

- 3. Execute the `/home/student/bin/bash-lab` script, and review the output content on `workstation`.

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade console-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish console-review
```

► Solution

Improve Command-line Productivity

Create a Bash script that can filter and get relevant information from different hosts.

Outcomes

- Create a Bash script and redirect its output to a file.
- Use loops to simplify your code.
- Filter the relevant content by using `grep` and regular expressions.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start console-review
```

Instructions

1. Create the executable `/home/student/bin/bash-lab` script file on the `workstation` machine. The initial content in the script must use the shebang interpreter directive.

- 1.1. On the `workstation` machine, create the `/home/student/bin/` directory if needed.

```
[student@workstation ~]$ mkdir -p /home/student/bin
```

- 1.2. Use the `vim` command to create and edit the `/home/student/bin/bash-lab` script file.

```
[student@workstation ~]$ vim ~/bin/bash-lab
```

- 1.3. Insert the following text and save the file.

```
#!/usr/bin/bash
```

- 1.4. Make your script file executable.

```
[student@workstation ~]$ chmod a+x ~/bin/bash-lab
```

2. Edit your newly created script file to store the following information from the `servera` and `serverb` machines on the `workstation` machine. The systems use SSH keys for authentication, and therefore you do not require a password. Store the output of the listed

commands from the following table in the `/home/student/output-servera` and `/home/student/output-serverb` files respectively on the `workstation` machine. Use the hash sign (`#`) for differentiating the output of the successive commands in the output file.

Command or file	Content requested
<code>hostname -f</code>	Store the entire output.
<code>echo "#####"</code>	Append the hash signs to differentiate the following command.
<code>lscpu</code>	Get only the lines that start with the CPU string.
<code>echo "#####"</code>	Append the hash signs to differentiate the following command.
<code>/etc/selinux/config</code>	Ignore empty lines. Also, ignore lines that start with the <code>#</code> character.
<code>echo "#####"</code>	Append the hash signs to differentiate the following command.
<code>/var/log/secure</code>	Get all "Failed password" entries.
<code>echo "#####"</code>	Append the hash signs to differentiate the following command.

Save the required information to the `output-servera` and `output-serverb` files in the `/home/student` directory on `workstation`.



Note You can use the `sudo` command without requiring a password on the `servera` and `serverb` hosts. Remember to use a loop to simplify your script. You can also use multiple `grep` commands that are concatenated with the use of the pipe character (`|`).

- 2.1. Use the `vim` command to open and edit the `/home/student/bin/bash-lab` script file.

```
[student@workstation ~]$ vim ~/bin/bash-lab
```

- 2.2. Append the following lines to the `/home/student/bin/bash-lab` script file. The number of hash signs is arbitrary.



Note The following output is an example of how you can achieve the requested script. In Bash scripting, you can take different approaches and obtain the same result.

```
#!/usr/bin/bash
USR='student'
OUT='/home/student/output'
#
for SRV in servera serverb; do
  ssh ${USR}@${SRV} "hostname -f" > ${OUT}-${SRV}
  echo "#####">>> ${OUT}-${SRV}
```

```
ssh ${USR}@${SRV} "lscpu | grep '^CPU'" >> ${OUT}-${SRV}
echo "#####" >> ${OUT}-${SRV}
ssh ${USR}@${SRV} "grep -v '^$' /etc/selinux/config|grep -v '^#'" >> ${OUT}-${SRV}
echo "#####" >> ${OUT}-${SRV}
ssh ${USR}@${SRV} "sudo grep 'Failed password' /var/log/secure" >> ${OUT}-${SRV}
echo "#####" >> ${OUT}-${SRV}
done
```

3. Execute the `/home/student/bin/bash-lab` script, and review the output content on workstation.

- 3.1. On workstation, execute the `/home/student/bin/bash-lab` script.

```
[student@workstation ~]$ bash-lab
```

- 3.2. Review the content of the `/home/student/output-servera` and `/home/student/output-serverb` files.

```
[student@workstation ~]$ cat /home/student/output-servera
servera.lab.example.com
#####
CPU op-mode(s):          32-bit, 64-bit
CPU(s):                  2
CPU family:               6
#####
SELINUX=enforcing
SELINUXTYPE=targeted
#####
Apr  1 05:42:07 servera sshd[1275]: Failed password for invalid user operator1
from 172.25.250.9 port 42460 ssh2
Apr  1 05:42:09 servera sshd[1277]: Failed password for invalid user sysadmin1
from 172.25.250.9 port 42462 ssh2
Apr  1 05:42:11 servera sshd[1279]: Failed password for invalid user manager1 from
172.25.250.9 port 42464 ssh2
#####
[student@workstation ~]$ cat /home/student/output-serverb
serverb.lab.example.com
#####
CPU op-mode(s):          32-bit, 64-bit
CPU(s):                  2
CPU family:               6
#####
SELINUX=enforcing
SELINUXTYPE=targeted
#####
Apr  1 05:42:14 serverb sshd[1252]: Failed password for invalid user operator1
from 172.25.250.9 port 53494 ssh2
Apr  1 05:42:17 serverb sshd[1257]: Failed password for invalid user sysadmin1
from 172.25.250.9 port 53496 ssh2
Apr  1 05:42:19 serverb sshd[1259]: Failed password for invalid user manager1 from
172.25.250.9 port 53498 ssh2
#####
```


Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade console-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish console-review
```

Summary

- Create and execute Bash scripts to accomplish administration tasks.
- Use loops to iterate through a list of items from the command line and in a shell script.
- Use conditional structures to incorporate decision-making into shell scripts.
- Search for text in log and configuration files by using regular expressions and the `grep` command.

Chapter 2

Schedule Future Tasks

Goal

Schedule tasks to execute at a specific time and date.

Objectives

- Set up a command to run once at a future time.
- Schedule commands to run on a repeating schedule with a user's `crontab` file.
- Schedule commands to run on a repeating schedule with the system `crontab` file and directories.
- Enable and disable `systemd` timers, and configure a timer that manages temporary files.

Sections

- Schedule a Deferred User Job (and Guided Exercise)
- Schedule Recurring User Jobs (and Guided Exercise)
- Schedule Recurring System Jobs (and Guided Exercise)
- Manage Temporary Files (and Guided Exercise)
- Schedule Future Tasks (Quiz)

Schedule a Deferred User Job

Objectives

- Set up a command to run once at a future time.

Describe Deferred User Tasks

Sometimes you might need to run one or more commands at a specific future time. An example is a user who schedules a long-running maintenance task to occur in the middle of the night. Another example is a system administrator who is working on a firewall configuration and queues a safety job to reset the firewall settings to a former working state in ten minutes' time. The system administrator then deactivates the job before it runs, unless the new firewall configuration worked.

These scheduled commands are called *tasks* or *jobs*, and the *deferred* term indicates that these tasks run in the future.

One available solution for Red Hat Enterprise Linux users to schedule deferred tasks is the `at` command, which is installed and enabled by default. The `at` package provides the `atd` system daemon and the `at` and `atq` commands to interact with the daemon.

Any user can queue jobs for the `atd` daemon by using the `at` command. The `atd` daemon provides 26 queues, identified from `a` to `z`, where jobs in alphabetically later queues get lower system priority (with higher *nice* values, as discussed in a later chapter).

Schedule Deferred User Tasks

Use the `at` *TIMESPEC* command to start entering a new job to schedule. The `at` command then reads from `STDIN` (your keyboard) to obtain the commands to run. When manually entering commands, complete the input by pressing `Ctrl+D` on an empty line. You can use input redirection from a script file for entering more complex commands. For example, use the `at now +5min < myscript` command to schedule the commands in `myscript` to start in 5 minutes, without needing to type the commands manually in a terminal window.

The `at` command *TIMESPEC* argument accepts natural time specifications to describe when a job should run. For example, specify a time as `02:00pm`, `15:59`, `midnight`, or even `teatime`, followed by an optional date or number of days in the future.

The *TIMESPEC* argument expects time and date specifications in that order. If you provide the date and not the time, then the time defaults to the current time. If you provide the time and not the date, then the date is considered to be matched, and the jobs run when the time next matches.

The following example shows a job schedule without providing the date. The `at` command schedules the job for today or tomorrow depending whether the time has passed.

```
[user@host ~]$ date
Wed May 18 21:01:18 CDT 2022
[user@host ~]$ at 21:03 < myscript
job 3 at Wed May 18 21:03:00 2022
[user@host ~]$ at 21:00 < myscript
job 4 at Thu May 19 21:00:00 2022
```

The man pages for the `at` command and other documentation sources use lowercase to write the natural time specifications. You can use lowercase, sentence case, or uppercase. Here are examples of time specifications that you can use:

- `now +5min`
- `teatime tomorrow` (teatime is 16:00)
- `noon +4 days`
- `5pm august 3 2021`

For other valid time specifications, refer to the local `timespec` document listed in the references.

Inspect and Manage Deferred User Jobs

For an overview of the pending jobs for the current user, use the `atq` or the `at -l` command.

```
[user@host ~]$ atq
28 Mon May 16 05:13:00 2022 a user
29 Tue May 17 16:00:00 2022 h user
30 Wed May 18 12:00:00 2022 a user
```

In the preceding output, every line represents a different scheduled future job. The following description applies to the first line of the output:

- **28** is the unique job number.
- **Mon May 16 05:13:00 2022** is the execution date and time for the scheduled job.
- **a** indicates that the job is scheduled with the default queue `a`.
- **user** is the owner of the job (and the user that the job runs as).



Important

Unprivileged users can view and manage only their own jobs. The `root` user can view and manage all jobs.

Use the `at -c JOBNUMBER` command to inspect the commands that run when the `atd` daemon executes a job. This command shows the job's environment, which is set from the user's environment when they created the job, and the command syntax to run.

Remove Jobs from Schedule

The `atrm JOBNUMBER` command removes a scheduled job. Remove the scheduled job when you no longer need it, for example, when a remote firewall configuration succeeded, and you do not need to reset it.



References

`at(1)` and `atd(8)` man pages

`/usr/share/doc/at/timespec`

► Guided Exercise

Schedule a Deferred User Job

Use the `at` command to schedule several commands to run at specified future times.

Outcomes

- Schedule a job to run at a specified future time.
- Inspect the commands that a scheduled job runs.
- Delete the scheduled jobs.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start scheduling-at
```

Instructions

- 1. From `workstation`, open an SSH session to `servera` as the `student` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
```

- 2. Schedule a job to run in two minutes from now. Save the output of the `date` command to the `/home/student/myjob.txt` file.

- 2.1. Pass the `date >> /home/student/myjob.txt` string as the input to the `at` command, so that the job runs in two minutes from now.

```
[student@servera ~]$ echo "date >> /home/student/myjob.txt" | at now +2min
warning: commands will be executed using /bin/sh
job 1 at Thu Feb 16 18:51:16 2023
```

- 2.2. List the scheduled jobs.

```
[student@servera ~]$ atq
1          Thu Feb 16 18:51:16 2023 a student
```

- 2.3. Monitor the deferred jobs queue in real time. After the `atd` daemon executes, it removes the job from the queue.

The command updates the output of the `atq` command every two seconds, by default. After the `atd` daemon removes the deferred job from the queue, press `Ctrl+C` to exit the `watch` command and return to the shell prompt.

```
[student@servera ~]$ watch atq
Every 2.0s: atq                  servera.lab.example.com: Thu Feb 16 17:58:50 2023

1          Thu Feb 16 18:51:16 2023 a student
```

- 2.4. Verify that the contents of the `/home/student/myjob.txt` file match the output of the `date` command.

The output matches the output of the `date` command, which confirms that the scheduled job executed successfully.

```
[student@servera ~]$ cat myjob.txt
Thu Feb 16 06:51:16 PM EDT 2023
```

- ▶ 3. Interactively schedule a job in the `g` queue that runs at `teatime` (16:00). The job should print the `It's teatime` message to the `/home/student/tea.txt` file. Append the new messages to the `/home/student/tea.txt` file.

```
[student@servera ~]$ at -q g teatime
warning: commands will be executed using /bin/sh
at> echo "It's teatime" >> /home/student/tea.txt
at> Ctrl+d
job 2 at Fri Feb 17 16:00:00 2023
```

- ▶ 4. Interactively schedule another job with the `b` queue that runs at `16:05`. The job should print `The cookies are good` message to the `/home/student/cookies.txt` file. Append the new messages to the `/home/student/cookies.txt` file.

```
[student@servera ~]$ at -q b 16:05
warning: commands will be executed using /bin/sh
at> echo "The cookies are good" >> /home/student/cookies.txt
at> Ctrl+d
job 3 at Fri Feb 17 16:05:00 2023
```

- ▶ 5. Inspect the commands in the pending jobs.

- 5.1. View the job numbers of the pending jobs.

Note the job numbers in the output, which might vary on your system. Use the job numbers from your system.

```
[student@servera ~]$ atq
2          Fri Feb 17 16:00:00 2023 g student
3          Fri Feb 17 16:05:00 2023 b student
```

- 5.2. View the commands in the pending job number 2. Replace the job number if it changed for you.

The job executes an `echo` command that appends the `It's teatime` message to the `/home/student/tea.txt` file.

```
[student@servera ~]$ at -c 2
...output omitted...
echo "It's teatime" >> /home/student/tea.txt
marcinDELIMITER1d7be6a7
```

- 5.3. View the commands in the pending job number 3. Replace the job number if it changed for you.

The job executes an `echo` command that appends the message `The cookies are good` to the `/home/student/cookies.txt` file.

```
[student@servera ~]$ at -c 3
...output omitted...
echo "The cookies are good" >> /home/student/cookies.txt
marcinDELIMITER44662c6f
```

- ▶ 6. View the job number of a job that runs at `teatime` (16:00), and remove it by using the `atrm` command.

```
[student@servera ~]$ atq
2      Fri Feb 17 16:00:00 2023 g student
3      Fri Feb 17 16:05:00 2023 b student
[student@servera ~]$ atrm 2
```

- ▶ 7. Verify that the scheduled job to run at `teatime` (16:00) no longer exists.
 - 7.1. View the list of pending jobs, and confirm that the scheduled job to run at `teatime` (16:00) no longer exists.

```
[student@servera ~]$ atq
3      Fri Feb 17 16:05:00 2023 b student
```

- 7.2. Return to the `workstation` machine as the `student` user.

```
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish scheduling-at
```

Schedule Recurring User Jobs

Objectives

- Schedule commands to run on a repeating schedule with a user's `crontab` file.

Describe Recurring User Jobs

Recurring jobs are scheduled to run repeatedly. Red Hat Enterprise Linux systems provide the `crond` daemon, which is enabled and started by default. The `crond` daemon reads multiple configuration files: one per user, and a set of system-wide files. Each user has a personal file that they edit with the `crontab -e` command. When executing recurring jobs, these configuration files provide detailed control to users and administrators. If the scheduled job is not written to use redirection, then the `crond` daemon emails any generated output or errors to the job owner.

Schedule Recurring User Jobs

Use the `crontab` command to manage scheduled jobs. The following list shows the commands that a local user can use to manage their jobs:

Examples Of The `crontab` Command

Command	Intended use
<code>crontab -l</code>	List the jobs for the current user.
<code>crontab -r</code>	Remove all jobs for the current user.
<code>crontab -e</code>	Edit jobs for the current user.
<code>crontab filename</code>	Remove all jobs, and replace them with jobs that are read from <i>filename</i> . This command uses <code>stdin</code> input when no file is specified.

A privileged user might use the `crontab` command `-u` option to manage jobs for another user. The `crontab` command is never used to manage system jobs, and using the `crontab` commands as the `root` user is not recommended due to the ability to exploit personal jobs that are configured to run as `root`. Configure such privileged jobs as described in the later section that describes recurring system jobs.

Describe User Job Format

The `crontab -e` command invokes the `vim` editor by default unless the `EDITOR` environment variable is set for another editor. Each job must use a unique line in the `crontab` file. Follow these recommendations for valid entries when writing recurring jobs:

- Empty lines for ease of reading
- Comments on lines that start with the number sign (`#`)
- Environment variables with a `NAME=value` format, which affects all lines after the line where they are declared

Standard variable settings include the `SHELL` variable, to declare the shell that is used for interpreting the remaining lines of the `crontab` file. The `MAILTO` variable determines who should receive the emailed output.

**Note**

The ability to send an email requires additional system configuration for a local mail server or an SMTP relay.

The fields in the `crontab` file appear in the following order:

- Minutes
- Hours
- Day of month
- Month
- Day of week
- Command

The command executes when the *Day of month* or *Day of week* fields use the same value other than the `*` character. For example, to run a command on the 11th day of every month, and every Friday at 12:15 (24-hour format), use the following job format:

```
15 12 11 * Fri command
```

The first five fields all use the same syntax rules:

- Use the `*` character to execute in every possible instance of the field.
- A number to specify the number of minutes or hours, a date, or a day of the week. For days of the week, 0 equals Sunday, 1 equals Monday, 2 equals Tuesday, and so on. 7 also equals Sunday.
- Use `x-y` for a range, which includes the `x` and `y` values.
- Use `x, y` for lists. Lists might include ranges as well, for example, `5, 10-13, 17` in the `Minutes` column, for a job to run at 5, 10, 11, 12, 13, and 17 minutes past the hour.
- The `*/x` indicates an interval of `x`; for example, `*/7` in the `Minutes` column runs a job every seven minutes.

Additionally, 3-letter English abbreviations are used for months or days of the week, for example, Jan, Feb, and Mon, Tue.

The last field contains the full command with options and arguments to execute with the default shell. If the command contains an unescaped percentage sign (`%`), then that percentage sign is treated as a newline character, and everything after the percentage sign passes to the command as `stdin` input.

Examples of Recurring User Jobs

The following job executes the `/usr/local/bin/yearly_backup` command at exactly 09:00 on 3 February, every year. February is represented as the number 2 in the example, because it is the second month of the year.

```
0 9 3 2 * /usr/local/bin/yearly_backup
```

The following job sends an email that contains the **Chime** word to the owner of this job every five minutes between and including 09:00 and 16:00, but only on each Friday in July.

```
* /5 9-16 * Jul 5 echo "Chime"
```

The preceding 9-16 range of hours means that the job timer starts at the ninth hour (09:00) and continues until the end of the sixteenth hour (16:59). The job starts executing at 09:00 with the last execution at 16:55, because five minutes after 16:55 is 17:00, which is beyond the given scope of hours.

If a range is specified for the hours instead of a single value, then all hours within the range will match. Therefore, with the hours of 9-16, this example matches every five minutes from 09:00 through 16:55.



Note

This example job sends the output as an email, because **crond** recognizes that the job allowed output to go to the **STDIO** channel without redirection. Because cron jobs run in a background environment without an output device (known as a *controlling terminal*), **crond** buffers the output and creates an email to send it to the specified user in the configuration. For system jobs, the email is sent to the root account.

The following job runs the `/usr/local/bin/daily_report` command every working day (Monday to Friday) two minutes before midnight.

```
58 23 * * 1-5 /usr/local/bin/daily_report
```

The following job executes the **mutt** command to send the **Checking in** mail message to the `developer@example.com` recipient every working day (Monday to Friday), at 9 AM.

```
0 9 * * 1-5 mutt -s "Checking in" developer@example.com % Hi there, just checking in.
```



References

`crond(8)`, `crontab(1)`, and `crontab(5)` man pages

► Guided Exercise

Schedule Recurring User Jobs

Schedule commands to run on a repeating schedule as a non-privileged user, with the `crontab` command.

Outcomes

- Schedule recurring jobs to run as a non-privileged user.
- Inspect the commands that a scheduled recurring job runs.
- Remove scheduled recurring jobs.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start scheduling-cron
```

Instructions

- 1. Log in to the `servera` machine as the `student` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
```

- 2. Schedule a recurring job as the `student` user that appends the current date and time to the `/home/student/my_first_cron_job.txt` file every two minutes. Use the `date` command to display the current date and time. The job must run only from one day before to one day after the current time. The job must not run on any other day.

- 2.1. Use the `date` command to display the current date and time. Note the day of the week, which you need for the next steps.

```
[student@workstation ~]$ date
Wed Mar 15 07:33:01 PM EDT 2023
[student@servera ~]$
```

**Note**

You can use the `date -d "last day" +%a` command to display the day before the current time, and the `date -d "next day" +%a` command to display the day after the current time.

```
[student@servera ~]$ date -d "last day" +%a
Tue
[student@servera ~]$ date -d "next day" +%a
Thu
```

- 2.2. Open the `crontab` file with the default text editor.

```
[student@servera ~]$ crontab -e
```

- 2.3. Insert the following line. Replace the range of days from one day before to one day after the current time:

```
*/2 * * * Tue-Thu /usr/bin/date >> /home/student/my_first_cron_job.txt
```

- 2.4. Press `ESC` and type `:wq` to save the changes and exit the editor. When the editor exits, you should see the following output:

```
...output omitted...
crontab: installing new crontab
[student@servera ~]$
```

- ▶ 3. Use the `crontab -l` command to list the scheduled recurring jobs. Inspect the command that you scheduled to run as a recurring job in the preceding step.

Verify that the job runs the `/usr/bin/date` command and appends its output to the `/home/student/my_first_cron_job.txt` file.

```
[student@servera ~]$ crontab -l
*/2 * * * Tue-Thu /usr/bin/date >> /home/student/my_first_cron_job.txt
```

- ▶ 4. Instruct your shell prompt to sleep until the `/home/student/my_first_cron_job.txt` file is created because of the successful execution of the recurring job that you scheduled. Wait for your shell prompt to return.

The `while` command uses `! test -f` to continue to run a loop, and sleeps for one second until the `my_first_cron_job.txt` file is created in the `/home/student` directory.

```
[student@servera ~]$ while ! test -f my_first_cron_job.txt; do sleep 1s; done
```

- ▶ 5. Verify that the contents of the `/home/student/my_first_cron_job.txt` file match the output of the `date` command.

```
[student@servera ~]$ cat my_first_cron_job.txt
Wed Mar 15 07:40:01 PM EDT 2023
```

- ▶ 6. Remove all the scheduled recurring jobs for the `student` user.

6.1. Remove all the scheduled recurring jobs for the `student` user.

```
[student@servera ~]$ crontab -r
```

6.2. Verify that no recurring jobs exist for the `student` user.

```
[student@servera ~]$ crontab -l
no crontab for student
```

6.3. Return to the `workstation` machine as the `student` user.

```
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish scheduling-cron
```

Schedule Recurring System Jobs

Objectives

- Schedule commands to run on a repeating schedule with the system crontab file and directories.

Recurring System Jobs

System administrators often need to run recurring jobs. It is best to run these jobs from system accounts rather than from user accounts. Schedule these jobs with system-wide crontab files instead of with the `crontab` command. Job entries in the system-wide crontab files are similar to the users' crontab entries. The system-wide crontab files have an extra field before the command field to specify the user that runs the command.

The `/etc/crontab` file has a syntax diagram in the comments.

```
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
```

The `/etc/crontab` file and other files in the `/etc/cron.d/` directory define the recurring system jobs. Always create custom crontab files in the `/etc/cron.d/` directory to schedule recurring system jobs. Place the custom crontab file in the `/etc/cron.d` directory to prevent a package update from overwriting the `/etc/crontab` file. Packages that require recurring system jobs place their crontab files in the `/etc/cron.d/` directory with the job entries. Administrators also use this location to group related jobs into a single file.

The crontab system also includes repositories for scripts to run every hour, day, week, and month. These repositories are placed in the `/etc/cron.hourly/`, `/etc/cron.daily/`, `/etc/cron.weekly/`, and `/etc/cron.monthly/` directories. These directories contain executable shell scripts, not crontab files.



Note

Use the `chmod +x script_name` command to make a script executable. A script must be executable to run.

Run Periodic Commands with Anacron

The `anacron` command uses the `run-parts` script to execute daily, weekly, and monthly jobs from the `/etc/anacrontab` configuration file.

The `/etc/anacrontab` file ensures that scheduled jobs always run and are not skipped accidentally because the system was turned off or hibernated. For example, when a system job that runs daily was not executed at a specified time because the system was rebooting, then the job is completed when the system becomes ready. A delay might occur before the job starts, if specified in the `Delay in minutes` parameter in the `/etc/anacrontab` file.

Files in the `/var/spool/anacron/` directory determine the daily, weekly, and monthly jobs. When the `crond` daemon starts a job from the `/etc/anacrontab` file, it updates the timestamps of those files. With this timestamp, you can determine the last time that the job executed. The syntax of the `/etc/anacrontab` file is different from the regular `crontab` configuration files. The `/etc/anacrontab` file contains four fields per line, as follows.

Period in days

Defines the interval in days for the job to run on a recurring schedule. This field accepts an integer or a macro value. For example, the macro `@daily` is equivalent to the `1` integer, which executes the job daily. Similarly, the macro `@weekly` is equivalent to the `7` integer, which executes the job weekly.

Delay in minutes

Defines the time that the `crond` daemon must wait before it starts the job.

Job identifier

Identifies the unique name of the job in the log messages.

Command

The command to be executed.

The `/etc/anacrontab` file also contains environment variable declarations with the `NAME=value` syntax. The `START_HOURS_RANGE` variable specifies the time interval for the jobs to run. Jobs do not start outside this range. When a job does not run within this time interval on a particular day, then the job must wait until the next day for execution.

Systemd Timer

The `systemd` timer unit activates another unit of a different type (such as a service), whose unit name matches the timer unit name. The timer unit allows timer-based activation of other units. The `systemd` timer unit logs timer events in system journals for easier debugging.

Sample Timer Unit

The `sysstat` package provides the `systemd` timer unit, called the `sysstat-collect.timer` service, to collect system statistics every 10 minutes. The following output shows the contents of the `/usr/lib/systemd/system/sysstat-collect.timer` configuration file.

```
...output omitted...
[Unit]
Description=Run system activity accounting tool every 10 minutes

[Timer]
OnCalendar=*:00/10
```

```
[Install]
WantedBy=sysstat.service
```

The `OnCalendar=*:00/10` option signifies that this timer unit activates the corresponding `sysstat-collect.service` unit every 10 minutes. You might specify more complex time intervals.

For example, a `2022-04-* 12:35,37,39:16` value against the `OnCalendar` option causes the timer unit to activate the corresponding service unit at the `12:35:16`, `12:37:16`, and `12:39:16` times, every day during April 2022. You might also specify relative timers with the `OnUnitActiveSec` option. For example, with the `OnUnitActiveSec=15min` option, the timer unit triggers the corresponding unit to start 15 minutes after the last time that the timer unit activated its corresponding unit.



Important

Do not modify any units in the configuration files under the `/usr/lib/systemd/system` directory, because the `systemd` unit overrides the configuration changes in that file. Create a copy of the configuration file in the `/etc/systemd/system` directory, and then modify the copied file to prevent any update to the provider package from overriding the changes. If two files exist with the same name in the `/usr/lib/systemd/system` and `/etc/systemd/system` directories, then the `systemd` timer unit parses the file in the `/etc/systemd/system` directory.

After you change the timer unit configuration file, use the `systemctl daemon-reload` command to ensure that the `systemd` timer unit loads the changes.

```
[root@host ~]# systemctl daemon-reload
```

After reloading the `systemd` daemon configuration, use the `systemctl` command to activate the timer unit.

```
[root@host ~]# systemctl enable --now <unitname>.timer
```



References

`crontab(5)`, `anacron(8)`, `anacrontab(5)`, `systemd.time(7)`, `systemd.timer(5)`, and `crond(8)` man pages

► Guided Exercise

Schedule Recurring System Jobs

Schedule commands to run on various schedules by adding configuration files to the system crontab directories.

Outcomes

- Schedule a recurring system job to count the number of active users.
- Update the systemd timer unit that gathers system activity data.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start scheduling-system
```

Instructions

- 1. Log in to the `servera` machine as the `student` user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

- 2. Schedule a recurring system job that generates a log message to indicate the number of active users in the system. This job must run daily and use the `w -h | wc -l` command to retrieve the number of active users in the system. Use the `logger` command to generate the log message of currently active users.

2.1. Create the `/etc/cron.daily/usercount` script file with the following content:

```
#!/bin/bash
USERCOUNT=$(w -h | wc -l)
logger "There are currently ${USERCOUNT} active users"
```

2.2. Make the script file executable.

```
[root@servera ~]# chmod +x /etc/cron.daily/usercount
```

- ▶ 3. Install the `sysstat` package. The timer unit must trigger the service unit every ten minutes to collect system activity data with the `/usr/lib64/sa/sa1` shell script. Change the timer unit configuration file to collect the system activity data every two minutes.

3.1. Install the `sysstat` package.

```
[root@servera ~]# dnf install sysstat
...output omitted...
Is this ok [y/N]: y
...output omitted...
Complete!
```

3.2. Copy the `/usr/lib/systemd/system/sysstat-collect.timer` file to the `/etc/systemd/system/sysstat-collect.timer` file.

```
[root@servera ~]# cp /usr/lib/systemd/system/sysstat-collect.timer \
/etc/systemd/system/sysstat-collect.timer
```

3.3. Edit the `/etc/systemd/system/sysstat-collect.timer` file for the timer unit to run every two minutes. Replace any occurrence of the `10 minutes` string with `2 minutes` throughout the unit configuration file, including the occurrences in the commented lines. Use the `vim /etc/systemd/system/sysstat-collect.timer` command to edit the configuration file.

From these changes, the `sysstat-collect.timer` unit triggers the `sysstat-collect.service` unit every two minutes, and collects the system activity data in a binary file in the `/var/log/sa` directory.

```
...output omitted...
# Activates activity collector every 2 minutes

[Unit]
Description=Run system activity accounting tool every 2 minutes

[Timer]
OnCalendar=*:00/2

[Install]
WantedBy=sysstat.service
```

3.4. Notify the `systemd` daemon of the changes.

```
[root@servera ~]# systemctl daemon-reload
```

3.5. Activate the `sysstat-collect.timer` unit.

```
[root@servera ~]# systemctl enable --now sysstat-collect.timer
...output omitted...
```

3.6. Wait until the binary file is created in the `/var/log/sa` directory.

The `while` command, `ls /var/log/sa | wc -l` returns `0` when the file does not exist, or returns `1` when the file exists. The `while` command pauses for one second when the file is not present. The `while` loop exits when the file is present.

```
[root@servera ~]# while [ $(ls /var/log/sa | wc -l) -eq 0 ]; \
do sleep 1s; done
```

- 3.7. Verify that the binary file in the `/var/log/sa` directory was modified within two minutes.

```
[root@servera ~]# ls -l /var/log/sa
total 4
-rw-r--r--. 1 root root 2540 Apr  5 04:08 sa05
[root@servera ~]# date
Tue Apr  5 04:08:29 AM EDT 2022
```

- 3.8. Return to the `workstation` machine as the `student` user.

```
[root@servera ~]# exit
logout
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish scheduling-system
```

Manage Temporary Files

Objectives

- Enable and disable `systemd` timers, and configure a timer that manages temporary files.

Manage Temporary Files

Most critical applications and services use temporary files and directories. Some applications and users use the `/tmp` directory to hold transient working data, whereas other applications use task-specific locations such as daemon- and user-specific volatile directories under `/run`, which exist only in memory. When the system reboots or loses power, memory-based file systems are self-cleaning.

Commonly, daemons and scripts operate correctly only when their expected temporary files and directories exist. Additionally, purging temporary files on persistent storage is necessary to prevent disk space issues or stale working data.

Red Hat Enterprise Linux includes the `systemd-tmpfiles` tool, which provides a structured and configurable method to manage temporary directories and files.

At system boot, one of the first `systemd` service units to launch is the `systemd-tmpfiles-setup` service. This service runs the `systemd-tmpfiles` command `--create --remove` option, which reads instructions from the `/usr/lib/tmpfiles.d/*.conf`, `/run/tmpfiles.d/*.conf`, and `/etc/tmpfiles.d/*.conf` configuration files. These configuration files list files and directories that the `systemd-tmpfiles-setup` service is instructed to create, delete, or secure with permissions.

Clean Temporary Files with a Systemd Timer

To prevent long-running systems from filling up their disks with stale data, a `systemd` timer unit called `systemd-tmpfiles-clean.timer` triggers at a regular interval the `systemd-tmpfiles-clean.service` unit, which executes the `systemd-tmpfiles --clean` command.

A `systemd` timer unit configuration has a `[Timer]` section to indicate how to start the service with the same name as the timer.

Use the following `systemctl` command to view the contents of the `systemd-tmpfiles-clean.timer` unit configuration file.

```
[user@host ~]$ systemctl cat systemd-tmpfiles-clean.timer
# /usr/lib/systemd/system/systemd-tmpfiles-clean.timer
# SPDX-License-Identifier: LGPL-2.1-or-later
#
# 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=Daily Cleanup of Temporary Directories
Documentation=man:tmpfiles.d(5) man:systemd-tmpfiles(8)
ConditionPathExists=!/etc/initrd-release

[Timer]
OnBootSec=15min
OnUnitActiveSec=1d
```

In the preceding configuration, the `OnBootSec=15min` parameter indicates that the `systemd-tmpfiles-clean.service` unit gets triggered 15 minutes after the system boots up. The `OnUnitActiveSec=1d` parameter indicates that any further trigger to the `systemd-tmpfiles-clean.service` unit happens 24 hours after the service unit was last activated.

Change the parameters in the `systemd-tmpfiles-clean.timer` unit configuration file to meet your requirements. For example, a `30min` value for the `OnUnitActiveSec` parameter triggers the `systemd-tmpfiles-clean.service` unit 30 minutes after the service unit is last activated. As a result, the `systemd-tmpfiles-clean.service` unit gets triggered every 30 minutes after the changes are recognized.

After you change the timer unit configuration file, use the `systemctl daemon-reload` command to ensure that the `systemd` daemon loads the new configuration.

```
[root@host ~]# systemctl daemon-reload
```

Clean Temporary Files Manually

The `systemd-tmpfiles --clean` command parses the same configuration files as the `systemd-tmpfiles --create` command, but instead of creating files and directories, it purges all files that were not accessed, changed, or modified more recently than the maximum age that is defined in the configuration file.

For detailed information about the format of the configuration files for the `systemd-tmpfiles` service, see the `tmpfiles.d(5)` man page. The syntax consists of the following columns: Type, Path, Mode, UID, GID, Age, and Argument. Type refers to the action for the `systemd-tmpfiles` service to take; for example, `d` to create a directory if it does not exist, or `Z` to recursively restore SELinux contexts, file permissions, and ownership.

The following command purges a configuration with explanations:

```
d /run/systemd/seats 0755 root root -
```

When you create files and directories, create the `/run/systemd/seats` directory if it does not exist, with the `root` user and the `root` group as owners, and with permissions of `rwxr-xr-x`. If this directory does exist, then take no action. The `systemd-tmpfiles` service does not purge this directory automatically.

```
D /home/student 0700 student student 1d
```

Create the `/home/student` directory if it does not exist. If it does exist, then remove all its contents. When the system runs the `systemd-tmpfiles --clean` command, it removes from the directory all files that you did not access, change, or modify for more than one day.

```
L /run/fstablink - root root - /etc/fstab
```

Create the `/run/fstablink` symbolic link, to point to the `/etc/fstab` directory. Never automatically purge this line.

Configuration File Precedence

The `systemd-tmpfiles-clean` service configuration files can exist in three places:

- `/etc/tmpfiles.d/*.conf`
- `/run/tmpfiles.d/*.conf`
- `/usr/lib/tmpfiles.d/*.conf`

Use the files in the `/etc/tmpfiles.d/` directory to configure custom temporary locations, and to override vendor-provided defaults. The files in the `/run/tmpfiles.d/` directory are volatile files, which normally daemons use to manage their own runtime temporary files. Relevant RPM packages provide the files in the `/usr/lib/tmpfiles.d/` directory; therefore do not edit these files.

If a file in the `/run/tmpfiles.d/` directory has the same file name as a file in the `/usr/lib/tmpfiles.d/` directory, then the service uses the file in the `/run/tmpfiles.d/` directory. If a file in the `/etc/tmpfiles.d/` directory has the same file name as a file in either the `/run/tmpfiles.d/` or the `/usr/lib/tmpfiles.d/` directories, then the service uses the file in the `/etc/tmpfiles.d/` directory.

Given these precedence rules, you can override vendor-provided settings by copying the relevant file to the `/etc/tmpfiles.d/` directory and then editing it. By using these configuration locations correctly, you can manage administrator-configured settings from a central configuration management system, and package updates do not overwrite your configured settings.



Note

When testing new or modified configurations, apply only the commands from a single configuration file at a time. Specify the name of the single configuration file on the `systemd-tmpfiles` command line.



References

`systemd-tmpfiles(8)`, `tmpfiles.d(5)`, `stat(1)`, `stat(2)`, and `systemd.timer(5)` man pages

► Guided Exercise

Manage Temporary Files

Configure `systemd-tmpfiles` to change how quickly it removes temporary files from the `/tmp` directory, and also to periodically purge files from another directory.

Outcomes

- Configure `systemd-tmpfiles` to remove unused temporary files from the `/tmp` directory.
- Configure `systemd-tmpfiles` to periodically purge files from another directory.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start scheduling-tmpfiles
```

Instructions

- 1. Log in to the `servera` system as the student user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

- 2. Configure the `systemd-tmpfiles` service to clean the `/tmp` directory of any unused files from the last five days. Ensure that a package update does not overwrite the configuration files.

- 2.1. Copy the `/usr/lib/tmpfiles.d/tmp.conf` file to the `/etc/tmpfiles.d` directory.

```
[root@servera ~]# cp /usr/lib/tmpfiles.d/tmp.conf \
/etc/tmpfiles.d/tmp.conf
```

- 2.2. Search for the configuration line in the `/etc/tmpfiles.d/tmp.conf` file that applies to the `/tmp` directory. Replace the existing age of the temporary files in that configuration line with the new age of 5 days. Remove from the file all the other lines, including the commented lines. You can use the `vim /etc/tmpfiles.d/tmp.conf` command to edit the configuration file.

In the configuration, the `q` type is the same as the `d` type, and instructs the `systemd-tmpfiles` service to create the `/tmp` directory if it does not exist. The directory's octal permissions must be set to `1777`. Both the owning user and group of the `/tmp` directory must be `root`. The `/tmp` directory must not contain the unused temporary files from the last five days.

The `/etc/tmpfiles.d/tmp.conf` file should appear as follows:

```
q /tmp 1777 root root 5d
```

2.3. Verify the `/etc/tmpfiles.d/tmp.conf` file configuration.

If the exit code is `0`, then the configuration settings are correct.

```
[root@servera ~]# systemd-tmpfiles --clean /etc/tmpfiles.d/tmp.conf
[root@servera ~]# echo $?
0
```

- ▶ 3. Add a new configuration that ensures that the `/run/momentary` directory exists, and that user and group ownership is set to the `root` user. The octal permissions for the directory must be `0700`. The configuration must purge from this directory any files that remain unused in the last 30 seconds.

3.1. Create the `/etc/tmpfiles.d/momentary.conf` file with the following content.

With the configuration, the `systemd-tmpfiles` service ensures that the `/run/momentary` directory exists and that its octal permissions are set to `0700`. The ownership of the `/run/momentary` directory must be the `root` user and group. The service purges from this directory any file if it remains unused for 30 seconds.

```
[root@servera ~]# vim /etc/tmpfiles.d/momentary.conf
d /run/momentary 0700 root root 30s
```

3.2. Verify the `/etc/tmpfiles.d/momentary.conf` file configuration. The command creates the `/run/momentary` directory if it does not exist.

Because the command does not return any errors, it confirms that the configuration settings are correct.

```
[root@servera ~]# systemd-tmpfiles --create \
/etc/tmpfiles.d/momentary.conf
```

3.3. Verify that the `systemd-tmpfiles` command creates the `/run/momentary` directory with the appropriate permissions, owner, and group owner.

The octal permission for the `/run/momentary` directory is set to `0700`, and the user and group ownership are set to `root`.

```
[root@servera ~]# ls -ld /run/momentary
drwx-----. 2 root root 40 Apr  4 06:35 /run/momentary
```

- ▶ 4. Verify that the `systemd-tmpfiles --clean` command removes from the `/run/momentary` directory any file that is unused in the last 30 seconds, based on the `systemd-tmpfiles` configuration for the directory.

- 4.1. Create the `/run/momentary/test` file.

```
[root@servera ~]# touch /run/momentary/test
```

- 4.2. Configure your shell prompt not to return for 30 seconds.

```
[root@servera ~]# sleep 30
```

- 4.3. After your shell prompt returns, clean stale files from the `/run/momentary` directory, based on the referenced rule in the `/etc/tmpfiles.d/momentary.conf` configuration file.

The command removes the `/run/momentary/test` file, because it remains unused for 30 seconds. This behavior is based on the referenced rule in the `/etc/tmpfiles.d/momentary.conf` configuration file.

```
[root@servera ~]# systemd-tmpfiles --clean \
/etc/tmpfiles.d/momentary.conf
```

- 4.4. Verify that the `/run/momentary/test` file does not exist.

```
[root@servera ~]# ls -l /run/momentary/test
ls: cannot access '/run/momentary/test': No such file or directory
```

- 4.5. Return to the `workstation` machine as the `student` user.

```
[root@servera ~]# exit
logout
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish scheduling-tempfiles
```

► Quiz

Schedule Future Tasks

Choose the correct answers to the following questions.

- 1. Which command displays all the user jobs that you scheduled to run as deferred jobs?
 - a. atq
 - b. atrm
 - c. at -c
 - d. at --display
- 2. Which command removes the deferred user job with the job number 5?
 - a. at -c 5
 - b. atrm 5
 - c. at 5
 - d. at --delete 5
- 3. Which command displays all the scheduled recurring user jobs for the currently logged-in user?
 - a. crontab -r
 - b. crontab -l
 - c. crontab -u
 - d. crontab -V
- 4. Which job format executes the `/usr/local/bin/daily_backup` command hourly from 9 AM to 6 PM on all days from Monday through Friday?
 - a. `00 * * * Mon-Fri /usr/local/bin/daily_backup`
 - b. `* */9 * * Mon-Fri /usr/local/bin/daily_backup`
 - c. `00 */18 * * * /usr/local/bin/daily_backup`
 - d. `00 09-18 * * Mon-Fri /usr/local/bin/daily_backup`
- 5. Which directory contains the shell scripts to run daily?
 - a. `/etc/cron.d`
 - b. `/etc/cron.hourly`
 - c. `/etc/cron.daily`
 - d. `/etc/cron.weekly`

- ▶ 6. Which configuration file defines the settings for the system jobs that run daily, weekly, and monthly?
 - a. `/etc/crontab`
 - b. `/etc/anacrontab`
 - c. `/etc/inittab`
 - d. `/etc/sysconfig/crond`

- ▶ 7. Which systemd unit regularly triggers the cleanup of temporary files?
 - a. `systemd-tmpfiles-clean.timer`
 - b. `systemd-tmpfiles-clean.service`
 - c. `dnf-makecache.timer`
 - d. `unbound-anchor.timer`

► Solution

Schedule Future Tasks

Choose the correct answers to the following questions.

- 1. Which command displays all the user jobs that you scheduled to run as deferred jobs?
 - a. `atq`
 - b. `atrm`
 - c. `at -c`
 - d. `at --display`
- 2. Which command removes the deferred user job with the job number 5?
 - a. `at -c 5`
 - b. `atrm 5`
 - c. `at 5`
 - d. `at --delete 5`
- 3. Which command displays all the scheduled recurring user jobs for the currently logged-in user?
 - a. `crontab -r`
 - b. `crontab -l`
 - c. `crontab -u`
 - d. `crontab -V`
- 4. Which job format executes the `/usr/local/bin/daily_backup` command hourly from 9 AM to 6 PM on all days from Monday through Friday?
 - a. `00 * * * Mon-Fri /usr/local/bin/daily_backup`
 - b. `* */9 * * Mon-Fri /usr/local/bin/daily_backup`
 - c. `00 */18 * * * /usr/local/bin/daily_backup`
 - d. `00 09-18 * * Mon-Fri /usr/local/bin/daily_backup`
- 5. Which directory contains the shell scripts to run daily?
 - a. `/etc/cron.d`
 - b. `/etc/cron.hourly`
 - c. `/etc/cron.daily`
 - d. `/etc/cron.weekly`

- ▶ 6. Which configuration file defines the settings for the system jobs that run daily, weekly, and monthly?
 - a. `/etc/crontab`
 - b. `/etc/anacrontab`
 - c. `/etc/inittab`
 - d. `/etc/sysconfig/crond`

- ▶ 7. Which systemd unit regularly triggers the cleanup of temporary files?
 - a. `systemd-tmpfiles-clean.timer`
 - b. `systemd-tmpfiles-clean.service`
 - c. `dnf-makecache.timer`
 - d. `unbound-anchor.timer`

Summary

- Deferred jobs or tasks are scheduled to run once in the future.
- Recurring user jobs execute the user's tasks on a repeating schedule.
- Recurring system jobs accomplish, on a repeating schedule, administrative tasks with system-wide impact.
- The `systemd` timer units can execute both the deferred and recurring jobs.

Chapter 3

Analyze and Store Logs

Goal

Locate and accurately interpret system event logs for troubleshooting purposes.

Objectives

- Describe the basic Red Hat Enterprise Linux logging architecture to record events.
- Interpret events in the relevant syslog files to troubleshoot problems or to review system status.
- Find and interpret entries in the system journal to troubleshoot problems or review system status.
- Configure the system journal to preserve the record of events when a server is rebooted.
- Maintain accurate time synchronization with Network Time Protocol (NTP) and configure the time zone to ensure correct time stamps for events that are recorded by the system journal and logs.

Sections

- Describe System Log Architecture (and Quiz)
- Review Syslog Files (and Guided Exercise)
- Review System Journal Entries (and Guided Exercise)
- Preserve the System Journal (and Guided Exercise)
- Maintain Accurate Time (and Guided Exercise)

Lab

- Analyze and Store Logs

Describe System Log Architecture

Objectives

- Describe the basic Red Hat Enterprise Linux logging architecture to record events.

System Logging

The operating system kernel and other processes record a log of events that happen when the system is running. These logs are used to audit the system and to troubleshoot problems. You can use text utilities such as the `less` and `tail` commands to inspect these logs.

Red Hat Enterprise Linux uses a standard logging system that is based on the syslog protocol to log the system messages. Many programs use the logging system to record events and to organize them into log files. The `systemd-journald` and `rsyslog` services handle the syslog messages in Red Hat Enterprise Linux 9.

The `systemd-journald` service is at the heart of the operating system event logging architecture. The `systemd-journald` service collects event messages from many sources:

- System kernel
- Output from the early stages of the boot process
- Standard output and standard error from daemons
- Syslog events

The `systemd-journald` service restructures the logs into a standard format and writes them into a structured, indexed system journal. By default, this journal is stored on a file system that does not persist across reboots.

The `rsyslog` service reads syslog messages that the `systemd-journald` service receives from the journal when they arrive. The `rsyslog` service then processes the syslog events, and records them to its log files or forwards them to other services according to its own configuration.

The `rsyslog` service sorts and writes syslog messages to the log files that do persist across reboots in the `/var/log` directory. The service also sorts the log messages to specific log files according to the type of program that sent each message and the priority of each syslog message.

In addition to syslog message files, the `/var/log` directory contains log files from other services on the system. The following table lists some useful files in the `/var/log` directory.

Selected System Log Files

Log file	Type of stored messages
<code>/var/log/messages</code>	Most syslog messages are logged here. Exceptions include messages about authentication and email processing, scheduled job execution, and purely debugging-related messages.
<code>/var/log/secure</code>	Syslog messages about security and authentication events.
<code>/var/log/maillog</code>	Syslog messages about the mail server.

Log file	Type of stored messages
/var/log/cron	Syslog messages about scheduled job execution.
/var/log/boot.log	Non-syslog console messages about system startup.

Some applications do not use the `syslog` service to manage their log messages. For example, the Apache Web Server saves log messages to files in a subdirectory of the `/var/log` directory.



References

`systemd-journald.service(8)`, `rsyslogd(8)`, and `rsyslog.conf(5)` man pages

For more information, refer to the *Troubleshooting Problems Using Log Files* section in the *Red Hat Enterprise Linux 9 Configuring Basic System Settings* guide at https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/configuring_basic_system_settings/index

► Quiz

Describe System Log Architecture

Choose the correct answer to the following questions:

- 1. Which log file stores most syslog messages, except for the ones about authentication, mail, scheduled jobs, and debugging?
 - a. /var/log/maillog
 - b. /var/log/boot.log
 - c. /var/log/messages
 - d. /var/log/secure
- 2. Which log file stores syslog messages about security and authentication operations in the system?
 - a. /var/log/maillog
 - b. /var/log/boot.log
 - c. /var/log/messages
 - d. /var/log/secure
- 3. Which service sorts and organizes syslog messages into files in the /var/log directory?
 - a. rsyslog
 - b. systemd-journald
 - c. auditd
 - d. tuned
- 4. Which directory accommodates the human-readable syslog files?
 - a. /sys/kernel/debug
 - b. /var/log/journal
 - c. /run/log/journal
 - d. /var/log
- 5. Which file stores syslog messages about the mail server?
 - a. /var/log/lastlog
 - b. /var/log/maillog
 - c. /var/log/tallylog
 - d. /var/log/boot.log

- ▶ 6. Which file stores syslog messages about scheduled jobs?
 - a. /var/log/cron
 - b. /var/log/tallylog
 - c. /var/log/spooler
 - d. /var/log/secure

- ▶ 7. Which file stores console messages about system startup?
 - a. /var/log/messages
 - b. /var/log/cron
 - c. /var/log/boot.log
 - d. /var/log/secure

► Solution

Describe System Log Architecture

Choose the correct answer to the following questions:

- 1. Which log file stores most syslog messages, except for the ones about authentication, mail, scheduled jobs, and debugging?
 - a. /var/log/maillog
 - b. /var/log/boot.log
 - c. /var/log/messages
 - d. /var/log/secure
- 2. Which log file stores syslog messages about security and authentication operations in the system?
 - a. /var/log/maillog
 - b. /var/log/boot.log
 - c. /var/log/messages
 - d. /var/log/secure
- 3. Which service sorts and organizes syslog messages into files in the /var/log directory?
 - a. rsyslog
 - b. systemd-journald
 - c. auditd
 - d. tuned
- 4. Which directory accommodates the human-readable syslog files?
 - a. /sys/kernel/debug
 - b. /var/log/journal
 - c. /run/log/journal
 - d. /var/log
- 5. Which file stores syslog messages about the mail server?
 - a. /var/log/lastlog
 - b. /var/log/maillog
 - c. /var/log/tallylog
 - d. /var/log/boot.log

- ▶ **6. Which file stores syslog messages about scheduled jobs?**
 - a. `/var/log/cron`
 - b. `/var/log/tallylog`
 - c. `/var/log/spooler`
 - d. `/var/log/secure`

- ▶ **7. Which file stores console messages about system startup?**
 - a. `/var/log/messages`
 - b. `/var/log/cron`
 - c. `/var/log/boot.log`
 - d. `/var/log/secure`

Review Syslog Files

Objectives

- Interpret events in relevant syslog files to troubleshoot problems or review system status.

Log Events to the System

Many programs use the syslog protocol to log events to the system. Each log message is categorized by facility (the subsystem that produces the message) and priority (the message's severity).

The following table lists the standard syslog facilities:

Overview of Syslog Facilities

Code	Facility	Facility description
0	kern	Kernel messages
1	user	User-level messages
2	mail	Mail system messages
3	daemon	System daemon messages
4	auth	Authentication and security messages
5	syslog	Internal syslog messages
6	lpr	Printer messages
7	news	Network news messages
8	uucp	UUCP protocol messages
9	cron	Clock daemon messages
10	authpriv	Non-system authorization messages
11	ftp	FTP protocol messages
16-23	local0 to local7	Custom local messages

The following table lists the standard syslog priorities in descending order:

Overview of Syslog Priorities

Code	Priority	Priority description
0	emerg	System is unusable

Code	Priority	Priority description
1	alert	Action must be taken immediately
2	crit	Critical condition
3	err	Non-critical error condition
4	warning	Warning condition
5	notice	Normal but significant event
6	info	Informational event
7	debug	Debugging-level message

The `rsyslog` service uses the facility and priority of log messages to determine how to handle them. Rules configure this facility and priority in the `/etc/rsyslog.conf` file and in any file in the `/etc/rsyslog.d` directory with the `.conf` extension. Software packages can add rules by installing an appropriate file in the `/etc/rsyslog.d` directory.

Each rule that controls how to sort syslog messages has a line in one of the configuration files. The left side of each line indicates the facility and priority of the syslog messages that the rule matches. The right side of each line indicates which file to save the log message in (or where else to deliver the message). An asterisk (`*`) is a wildcard that matches all values.

For example, the following line in the `/etc/rsyslog.conf` file would record messages that are sent to the `authpriv` facility at any priority to the `/var/log/secure` file:

```
authpriv.* /var/log/secure
```

Sometimes, log messages match more than one rule in the `rsyslog.conf` file. In such cases, one message is stored in more than one log file. The `none` keyword in the priority field indicates that no messages for the indicated facility are stored in the given file, to limit stored messages.

Instead of being logged to a file, syslog messages can also be printed to the terminals of all logged-in users. The `rsyslog.conf` file has a setting to print all the syslog messages with the `emerg` priority to the terminals of all logged-in users.

Sample Rules of the rsyslog Service

```
#### 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.
# Don't 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

# Save boot messages also to boot.log
local7.*                              /var/log/boot.log
```

Note

The syslog subsystem has many more features beyond the scope of this course. To explore further, refer to the `rsyslog.conf(5)` man page and the extensive HTML documentation at `/usr/share/doc/rsyslog/html/index.html` that the `rsyslog-doc` package provides.

Log File Rotation

The `logrotate` command rotates log files to prevent them from taking too much space in the `/var/log` directory. When a log file is rotated, it is renamed with an extension that indicates the rotation date. For example, the previous `/var/log/messages` file is renamed to the `/var/log/messages-20220320` file when it is rotated on 2022-03-20. After the previous log file rotates, it creates a log file and notifies the service that wrote the log file.

After rotations during typically four weeks, the earliest log file is discarded to free disk space. A scheduled job runs the `logrotate` command daily to see the rotation requirement of any log files. Most log files rotate weekly; the `logrotate` command rotates some log files faster, or more slowly, or when they reach a specific size.

Analyze a Syslog Entry

Log messages start with the earliest message at the start and the latest message at the end of the log file. The `rsyslog` service uses a standard format for recording entries in log files. The following example explains the anatomy of a log message in the `/var/log/secure` log file.

```
Mar 20 20:11:48 localhost sshd[1433]: Failed password for student from 172.25.0.10
port 59344 ssh2
```

- **Mar 20 20:11:48** : Records the time stamp of the log entry.
- **localhost** : The host that sends the log message.
- **sshd[1433]** : The program or process name and PID number that sent the log message.
- **Failed password for ...** : The message that was sent.

Monitor Log Events

Monitoring log files for events is helpful to reproduce issues. The `tail -f /path/to/file` command outputs the last ten lines of the specified file and continues to output newly written lines in the file.

For example, to monitor for failed login attempts, run the `tail` command in one terminal, and then run in another terminal the `ssh` command as the `root` user while a user tries to log in to the system.

In the first terminal, run the `tail` command:

```
[root@host ~]# tail -f /var/log/secure
```

In the second terminal, run the `ssh` command:

```
[root@host ~]# ssh root@hosta
root@hosta's password: redhat
...output omitted...
[root@hostA ~]#
```

The log messages are visible in the first terminal.

```
...output omitted...
Mar 20 09:01:13 host sshd[2712]: Accepted password for root from 172.25.254.254
port 56801 ssh2
Mar 20 09:01:13 host sshd[2712]: pam_unix(sshd:session): session opened for user
root by (uid=0)
```

Send Syslog Messages Manually

The `logger` command sends messages to the `rsyslog` service. By default, the `logger` command sends the message to the user type with the `notice` priority (`user.notice`) unless specified otherwise with the `-p` option. It is helpful to test any change to the `rsyslog` service configuration.

To send a message to the `rsyslog` service to be recorded in the `/var/log/boot.log` log file, execute the following `logger` command:

```
[root@host ~]# logger -p local7.notice "Log entry created on host"
```



References

logger(1), tail(1), rsyslog.conf(5), and logrotate(8) man pages

rsyslog Manual

- `/usr/share/doc/rsyslog/html/index.html` provided by the `rsyslog-doc` package

For further information, refer to *Troubleshooting Problems Using Log Files* at https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/configuring_basic_system_settings/assembly_troubleshooting-problems-using-log-files_configuring-basic-system-settings

► Guided Exercise

Review Syslog Files

Reconfigure the `rsyslog` service to write specific log messages to a new file.

Outcomes

- Configure the `rsyslog` service to write all log messages with the `debug` priority to the `/var/log/messages-debug` log file.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start logs-syslog
```

Instructions

- 1. Log in to the `servera` machine as the `student` user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

- 2. Configure the `rsyslog` service on the `servera` machine to log all messages with the `debug` or higher priority, for any service to the new `/var/log/messages-debug` log file by changing the `/etc/rsyslog.d/debug.conf` configuration file.

- 2.1. Create the `/etc/rsyslog.d/debug.conf` file with the necessary entries to redirect all log messages with the `debug` or higher priority to the `/var/log/messages-debug` log file.

```
*.debug /var/log/messages-debug
```

This configuration line logs syslog messages with any facility and with the `debug` or higher priority level:

- The wildcard (*) in the facility field of the configuration line indicates any facility of log messages.
- The `rsyslog` service writes the matching messages to the `/var/log/messages-debug` log file.

- 2.2. Restart the `rsyslog` service.

```
[root@servera ~]# systemctl restart rsyslog
```

- ▶ 3. Verify that all the log messages with the debug priority appear in the `/var/log/messages-debug` log file.

3.1. Generate a log message with the user type and the debug priority.

```
[root@servera ~]# logger -p user.debug "Debug Message Test"
```

3.2. View the last ten log messages from the `/var/log/messages-debug` log file, and verify that you see the `Debug Message Test` message among the other log messages.

```
[root@servera ~]# tail /var/log/messages-debug
Feb 13 18:22:38 servera systemd[1]: Stopping System Logging Service...
Feb 13 18:22:38 servera rsyslogd[25176]: [origin software="rsyslogd"
swVersion="8.37.0-9.el8" x-pid="25176" x-info="http://www.rsyslog.com"] exiting
on signal 15.
Feb 13 18:22:38 servera systemd[1]: Stopped System Logging Service.
Feb 13 18:22:38 servera systemd[1]: Starting System Logging Service...
Feb 13 18:22:38 servera rsyslogd[25410]: environment variable TZ is not set, auto
correcting this to TZ=/etc/localtime [v8.37.0-9.el8 try http://www.rsyslog.com/
e/2442 ]
Feb 13 18:22:38 servera systemd[1]: Started System Logging Service.
Feb 13 18:22:38 servera rsyslogd[25410]: [origin software="rsyslogd"
swVersion="8.37.0-9.el8" x-pid="25410" x-info="http://www.rsyslog.com"] start
Feb 13 18:27:58 servera root[25416]: Debug Message Test
```

3.3. Return to the workstation system as the student user.

```
[root@servera ~]# exit
logout
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the workstation machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish logs-syslog
```

Review System Journal Entries

Objectives

- Find and interpret entries in the system journal to troubleshoot problems or review system status.

Find Events on the System Journal

The `systemd-journald` service stores logging data in a structured, indexed binary file called a *journal*. This data includes extra information about the log event. For example, for `syslog` events, this information includes the priority of the original message and the *facility*, which is a value that the `syslog` service assigns to track the process that originated a message.



Important

In Red Hat Enterprise Linux, the memory-based `/run/log` directory holds the system journal by default. The contents of the `/run/log` directory are lost when the system is shut down. You can change the `journald` directory to a persistent location, which is discussed later in this chapter.

To retrieve log messages from the journal, use the `journalctl` command. You can use the `journalctl` command to view all messages in the journal, or to search for specific events based on options and criteria. If you run the command as `root`, then you have full access to the journal. Although regular users can also use the `journalctl` command, the system restricts them from seeing certain messages.

```
[root@host ~]# journalctl
...output omitted...
Mar 15 04:42:16 host.lab.example.com systemd[2127]: Listening on PipeWire
Multimedia System Socket.
Mar 15 04:42:16 host.lab.example.com systemd[2127]: Starting Create User's
Volatile Files and Directories...
Mar 15 04:42:16 host.lab.example.com systemd[2127]: Listening on D-Bus User
Message Bus Socket.
Mar 15 04:42:16 host.lab.example.com systemd[2127]: Reached target Sockets.
Mar 15 04:42:16 host.lab.example.com systemd[2127]: Finished Create User's
Volatile Files and Directories.
Mar 15 04:42:16 host.lab.example.com systemd[2127]: Reached target Basic System.
Mar 15 04:42:16 host.lab.example.com systemd[1]: Started User Manager for UID 0.
Mar 15 04:42:16 host.lab.example.com systemd[2127]: Reached target Main User
Target.
Mar 15 04:42:16 host.lab.example.com systemd[2127]: Startup finished in 90ms.
Mar 15 04:42:16 host.lab.example.com systemd[1]: Started Session 6 of User root.
Mar 15 04:42:16 host.lab.example.com sshd[2110]: pam_unix(sshd:session): session
opened for user root(uid=0) by (uid=0)
```



```
Mar 15 04:42:17 host.lab.example.com systemd[1]: Starting Hostname Service...
Mar 15 04:42:17 host.lab.example.com systemd[1]: Started Hostname Service.
lines 1951-2000/2000 (END) q
```

The `journalctl` command highlights important log messages; messages with the `notice` or `warning` priority are in bold text, whereas messages with the `error` priority or higher are in red text.

The key to successful use of the journal for troubleshooting and auditing is to limit journal searches to show only relevant output.

By default, the `journalctl` command `-n` option shows the last 10 log entries. You can adjust the number of log entries with an optional argument that specifies how many log entries to display. For example, to review the last five log entries, you can run the following `journalctl` command:

```
[root@host ~]# journalctl -n 5
Mar 15 04:42:17 host.lab.example.com systemd[1]: Started Hostname Service.
Mar 15 04:42:47 host.lab.example.com systemd[1]: systemd-hostnamed.service:
Deactivated successfully.
Mar 15 04:47:33 host.lab.example.com systemd[2127]: Created slice User Background
Tasks Slice.
Mar 15 04:47:33 host.lab.example.com systemd[2127]: Starting Cleanup of User's
Temporary Files and Directories...
Mar 15 04:47:33 host.lab.example.com systemd[2127]: Finished Cleanup of User's
Temporary Files and Directories.
```

Similar to the `tail` command, the `journalctl` command `-f` option outputs the last 10 lines of the system journal and continues to output new journal entries when the journal appends them. To exit the `journalctl` command `-f` option, use the `Ctrl+C` key combination.

```
[root@host ~]# journalctl -f
Mar 15 04:47:33 host.lab.example.com systemd[2127]: Finished Cleanup of User's
Temporary Files and Directories.
Mar 15 05:01:01 host.lab.example.com CROND[2197]: (root) CMD (run-parts /etc/
cron.hourly)
Mar 15 05:01:01 host.lab.example.com run-parts[2200]: (/etc/cron.hourly) starting
0anacron
Mar 15 05:01:01 host.lab.example.com anacron[2208]: Anacron started on 2022-03-15
Mar 15 05:01:01 host.lab.example.com anacron[2208]: Will run job `cron.daily' in
29 min.
Mar 15 05:01:01 host.lab.example.com anacron[2208]: Will run job `cron.weekly' in
49 min.
Mar 15 05:01:01 host.lab.example.com anacron[2208]: Will run job `cron.monthly' in
69 min.
Mar 15 05:01:01 host.lab.example.com anacron[2208]: Jobs will be executed
sequentially
Mar 15 05:01:01 host.lab.example.com run-parts[2210]: (/etc/cron.hourly) finished
0anacron
Mar 15 05:01:01 host.lab.example.com CROND[2196]: (root) CMDEND (run-parts /etc/
cron.hourly)
^C
[root@host ~]#
```

To help to troubleshoot problems, you can filter the output of the journal by the priority of the journal entries. The `journalctl` command `-p` option shows the journal entries with a specified priority level (by name or by number) or higher. The `journalctl` command processes the `debug`, `info`, `notice`, `warning`, `err`, `crit`, `alert`, and `emerg` priority levels, in ascending priority order.

As an example, run the following `journalctl` command to list journal entries with the `err` priority or higher:

```
[root@host ~]# journalctl -p err
Mar 15 04:22:00 host.lab.example.com pipewire-pulse[1640]: pw.conf: execvp error
'pactl': No such file or direct
Mar 15 04:22:17 host.lab.example.com kernel: Detected CPU family 6 model 13
stepping 3
Mar 15 04:22:17 host.lab.example.com kernel: Warning: Intel Processor - this
hardware has not undergone testing by Red Hat and might not be certif>
Mar 15 04:22:20 host.lab.example.com smartd[669]: DEVICSCAN failed: glob(3)
aborted matching pattern /dev/discs/disc*
Mar 15 04:22:20 host.lab.example.com smartd[669]: In the system's table of devices
NO devices found to scan
```

You can show messages for a specified systemd unit by using the `journalctl` command `-u` option and the unit name.

```
[root@host ~]# journalctl -u sshd.service
May 15 04:30:18 host.lab.example.com systemd[1]: Starting OpenSSH server daemon...
May 15 04:30:18 host.lab.example.com sshd[1142]: Server listening on 0.0.0.0 port
22.
May 15 04:30:18 host.lab.example.com sshd[1142]: Server listening on :: port 22.
May 15 04:30:18 host.lab.example.com systemd[1]: Started OpenSSH server daemon.
May 15 04:32:03 host.lab.example.com sshd[1796]: Accepted publickey for user1 from
172.25.250.254 port 43876 ssh2: RSA SHA256:1UGy...>
May 15 04:32:03 host.lab.example.com sshd[1796]: pam_unix(sshd:session): session
opened for user user1(uid=1000) by (uid=0)
May 15 04:32:26 host.lab.example.com sshd[1866]: Accepted publickey for user2
from ::1 port 36088 ssh2: RSA SHA256:M8ik...
May 15 04:32:26 host.lab.example.com sshd[1866]: pam_unix(sshd:session): session
opened for user user2(uid=1001) by (uid=0)
lines 1-8/8 (END) q
```

When looking for specific events, you can limit the output to a specific time frame. To limit the output to a specific time range, the `journalctl` command has the `--since` option and the `--until` option. Both options take a time argument in the `"YYYY-MM-DD hh:mm:ss"` format (the double quotation marks are required to preserve the space in the option).

The `journalctl` command assumes that the day starts at 00:00:00 when you omit the time argument. The command assumes the current day when you omit the day argument. Both options take `yesterday`, `today`, and `tomorrow` as valid arguments in addition to the date and time field.

As an example, run the following `journalctl` command to list all journal entries from today's records:

```
[root@host ~]# journalctl --since today
...output omitted...
Mar 15 05:04:20 host.lab.example.com systemd[1]: Started Session 8 of User
student.
Mar 15 05:04:20 host.lab.example.com sshd[2255]: pam_unix(sshd:session): session
opened for user student(uid=1000) by (uid=0)
Mar 15 05:04:20 host.lab.example.com systemd[1]: Starting Hostname Service...
Mar 15 05:04:20 host.lab.example.com systemd[1]: Started Hostname Service.
Mar 15 05:04:50 host.lab.example.com systemd[1]: systemd-hostnamed.service:
Deactivated successfully.
Mar 15 05:06:33 host.lab.example.com systemd[2261]: Starting Mark boot as
successful...
Mar 15 05:06:33 host.lab.example.com systemd[2261]: Finished Mark boot as
successful.
lines 1996-2043/2043 (END) q
```

Run the following `journalctl` command to list all journal entries from 2022-03-11 20:30:00 to 2022-03-14 10:00:00:

```
[root@host ~]# journalctl --since "2022-03-11 20:30" --until "2022-03-14 10:00"
...output omitted...
```

You can also specify all entries since a relative time to the present. For example, to specify all entries in the last hour, you can use the following command:

```
[root@host ~]# journalctl --since "-1 hour"
...output omitted...
```



Note

You can use other, more sophisticated time specifications with the `--since` and `--until` options. For some examples, see the `systemd.time(7)` man page.

In addition to the visible content of the journal, you can view additional log entries if you turn on the verbose output. You can use any displayed extra field to filter the output of a journal query. The verbose output is useful to reduce the output of complex searches for certain events in the journal.

```
[root@host ~]# journalctl -o verbose
Tue 2022-03-15 05:10:32.625470 EDT [s=e7623387430b4c14b2c71917db58e0ee;i...]
  _BOOT_ID=beaadd6e5c5448e393ce716cd76229d4
  _MACHINE_ID=4ec03abd2f7b40118b1b357f479b3112
  PRIORITY=6
  SYSLOG_FACILITY=3
  SYSLOG_IDENTIFIER=systemd
  _UID=0
  _GID=0
  _TRANSPORT=journal
  _CAP_EFFECTIVE=1fffffffff
  TID=1
  CODE_FILE=src/core/job.c
```

```

CODE_LINE=744
CODE_FUNC=job_emit_done_message
JOB_RESULT=done
_PID=1
_COMM=systemd
_EXE=/usr/lib/systemd/systemd
_SYSTEMD_CGROUP=/init.scope
_SYSTEMD_UNIT=init.scope
_SYSTEMD_SLICE=-.slice
JOB_TYPE=stop
MESSAGE_ID=9d1aaa27d60140bd96365438aad20286
_HOSTNAME=host.lab.example.com
_CMDLINE=/usr/lib/systemd/systemd --switched-root --system --deserialize 31
_SELINUX_CONTEXT=system_u:system_r:init_t:s0
UNIT=user-1000.slice
MESSAGE=Removed slice User Slice of UID 1000.
INVOCATION_ID=0e5efc1b4a6d41198f0cf02116ca8aa8
JOB_ID=3220
_SOURCE_REALTIME_TIMESTAMP=1647335432625470
lines 46560-46607/46607 (END) q

```

The following list shows some fields of the system journal that you can use to search for relevant lines to a particular process or event:

- `_COMM` is the command name.
- `_EXE` is the path to the executable file for the process.
- `_PID` is the PID of the process.
- `_UID` is the UID of the user that runs the process.
- `_SYSTEMD_UNIT` is the `systemd` unit that started the process.

You can combine multiple system journal fields to form a granular search query with the `journalctl` command. For example, the following `journalctl` command shows all related journal entries to the `sshd.service` `systemd` unit from a process with PID 2110.

```

[root@host ~]# journalctl _SYSTEMD_UNIT=sshd.service _PID=2110
Mar 15 04:42:16 host.lab.example.com sshd[2110]: Accepted
publickey for root from 172.25.250.254 port 46224 ssh2: RSA
SHA256:1UGybTe52L2jzEJa1HLVKn9QUCKrTv3ZzxnmJol1Fro
Mar 15 04:42:16 host.lab.example.com sshd[2110]: pam_unix(sshd:session): session
opened for user root(uid=0) by (uid=0)

```



Note

For a list of journal fields, consult the `systemd.journal-fields(7)` man page.



References

`journalctl(1)`, `systemd.journal-fields(7)`, and `systemd.time(7)` man pages

For more information refer to the *Troubleshooting Problems Using Log Files* section in the *Red Hat Enterprise Linux 9 Configuring Basic System Settings* guide at https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/configuring_basic_system_settings/index#troubleshooting-problems-using-log-files_getting-started-with-system-administration

► Guided Exercise

Review System Journal Entries

Search the system journal for entries to record events that match specific criteria.

Outcomes

- Search the system journal for entries to record events based on different criteria.

Before You Begin

As the **student** user on the **workstation** machine, use the **lab** command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start logs-systemd
```

Instructions

- 1. From the **workstation** machine, open an SSH session to the **servera** machine as the **student** user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
```

- 2. Use the **journalctl** command **_PID=1** option to display only log events that originate from the **systemd** PID 1 process on the **servera** machine. To quit from the **journalctl** command, press **q**. The following output is an example and might differ on your system:

```
[student@servera ~]$ journalctl _PID=1
Mar 15 04:21:14 localhost systemd[1]: Finished Load Kernel Modules.
Mar 15 04:21:14 localhost systemd[1]: Finished Setup Virtual Console.
Mar 15 04:21:14 localhost systemd[1]: dracut ask for additional cmdline parameters
was skipped because all trigger condition checks failed.
Mar 15 04:21:14 localhost systemd[1]: Starting dracut cmdline hook...
Mar 15 04:21:14 localhost systemd[1]: Starting Apply Kernel Variables...
lines 1-5 q
[student@servera ~]$
```

- 3. Use the **journalctl** command **_UID=81** option to display all log events that originated from a system service with a UID of 81 on the **servera** machine.

```
[student@servera ~]$ journalctl _UID=81
Mar 15 04:21:17 servera.lab.example.com dbus-broker-lau[727]: Ready
```

- 4. Use the `journalctl` command `-p warning` option to display log events with a warning or higher priority on the `servera` machine.

```
[student@servera ~]$ journalctl -p warning
Mar 15 04:21:14 localhost kernel: wait_for_initramfs() called before
rootfs_initcalls
Mar 15 04:21:14 localhost kernel: ACPI: PRMT not present
Mar 15 04:21:14 localhost kernel: acpi PNP0A03:00: fail to add MMCONFIG
information, can't access extended PCI configuration space under this bridge.
Mar 15 04:21:14 localhost kernel: device-mapper: core: CONFIG_IMA_DISABLE_HTABLE
is disabled. Duplicate IMA measurements will not be recorded in the IMA log.
...output omitted...
Mar 15 04:21:18 servera.lab.example.com NetworkManager[769]: <warn>
[1647332478.5504] device (eth0): mtu: failure to set IPv6 MTU
Mar 15 04:21:27 servera.lab.example.com chronyd[751]: System clock wrong by
-0.919695 seconds
Mar 15 04:22:34 servera.lab.example.com chronyd[751]: System clock wrong by
0.772805 seconds
Mar 15 05:41:11 servera.lab.example.com sshd[1104]: error:
kex_exchange_identification: Connection closed by remote host
lines 1-19/19 (END) q
[student@servera ~]$
```

- 5. Display all recorded log events in the past 10 minutes from the current time on the `servera` machine.

```
[student@servera ~]$ journalctl --since "-10min"
Mar 15 05:40:01 servera.lab.example.com anacron[1092]: Job `cron.weekly' started
Mar 15 05:40:01 servera.lab.example.com anacron[1092]: Job `cron.weekly'
terminated
Mar 15 05:41:11 servera.lab.example.com sshd[1104]: error:
kex_exchange_identification: Connection closed by remote host
Mar 15 05:41:11 servera.lab.example.com sshd[1104]: Connection closed by
172.25.250.9 port 45370
Mar 15 05:41:14 servera.lab.example.com sshd[1105]: Accepted publickey for student
from 172.25.250.9 port 45372 ssh2: RSA SHA256:M8ikhcEDm2tQ95Z0o7ZvufqEixCFct
+wowZLNzNlBT0
Mar 15 05:41:14 servera.lab.example.com systemd[1]: Created slice User Slice of
UID 1000.
Mar 15 05:41:14 servera.lab.example.com systemd[1]: Starting User Runtime
Directory /run/user/1000...
Mar 15 05:41:14 servera.lab.example.com systemd-logind[739]: New session 1 of user
student.
Mar 15 05:41:14 servera.lab.example.com systemd[1]: Finished User Runtime
Directory /run/user/1000.
Mar 15 05:41:14 servera.lab.example.com systemd[1]: Starting User Manager for UID
1000...
...output omitted...
Mar 15 05:44:56 servera.lab.example.com systemd[1109]: Stopped target Sockets.
Mar 15 05:44:56 servera.lab.example.com systemd[1109]: Stopped target Timers.
```

```
Mar 15 05:44:56 servera.lab.example.com systemd[1109]: Stopped Mark boot as
successful after the user session has run 2 minutes.
Mar 15 05:44:56 servera.lab.example.com systemd[1109]: Stopped Daily Cleanup of
User's Temporary Directories.
lines 1-48 q
[student@servera ~]$
```

- ▶ 6. Use the `journalctl` command `--since` and `_SYSTEMD_UNIT="sshd.service"` options to display all the recorded log events that originated from the `sshd` service since 09:00:00 this morning on the `servera` machine.

**Note**

Online classrooms typically run on the UTC time zone. To obtain results that start at 9:00 AM in your local time zone, adjust your `--since` value by the amount of your offset from UTC. Alternatively, ignore the local time and use a value of 9:00 to locate journal entries that occurred since 9:00 for the `servera` time zone.

```
[student@servera ~]$ journalctl --since 9:00:00 _SYSTEMD_UNIT="sshd.service"
Mar 15 09:41:14 servera.lab.example.com sshd[1105]: Accepted publickey for student
from 172.25.250.9 port 45372 ssh2: RSA SHA256:M8ikhcEDm2tQ95Z0o7ZvufqEixCFct
+wowZLNzNlBT0
Mar 15 09:41:15 servera.lab.example.com sshd[1105]: pam_unix(sshd:session):
session opened for user student(uid=1000) by (uid=0)
Mar 15 09:44:56 servera.lab.example.com sshd[1156]: Accepted publickey for student
from 172.25.250.9 port 45374 ssh2: RSA SHA256:M8ikhcEDm2tQ95Z0o7ZvufqEixCFct
+wowZLNzNlBT0
Mar 15 09:44:56 servera.lab.example.com sshd[1156]: pam_unix(sshd:session):
session opened for user student(uid=1000) by (uid=0)
```

- ▶ 7. Return to the workstation system as the student user.

```
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish logs-systemd
```


Preserve the System Journal

Objectives

- Configure the system journal to preserve the record of events when a server is rebooted.

System Journal Storage

By default, Red Hat Enterprise Linux 9 stores the system journal in the `/run/log` directory, and the system clears the system journal after a reboot. You can change the configuration settings of the `systemd-journald` service in the `/etc/systemd/journald.conf` file so that the journals persist across a reboot.

The `Storage` parameter in the `/etc/systemd/journald.conf` file defines whether to store system journals in a volatile manner or persistently across a reboot. Set this parameter to `persistent`, `volatile`, `auto`, or `none`, as follows:

- **persistent**: Stores journals in the `/var/log/journal` directory, which persists across reboots. If the `/var/log/journal` directory does not exist, then the `systemd-journald` service creates it.
- **volatile**: Stores journals in the volatile `/run/log/journal` directory. Because the `/run` file system is temporary and exists only in the runtime memory, the data in it, including system journals, does not persist across a reboot.
- **auto**: If the `/var/log/journal` directory exists, then the `systemd-journald` service uses persistent storage; otherwise it uses volatile storage. This action is the default if you do not set the `Storage` parameter.
- **none**: Do not use any storage. The system drops all logs, but you can still forward the logs.

The advantage of persistent system journals is that the historical data is available immediately at boot. However, even with a persistent journal, the system does not keep all data forever. The journal has a built-in log rotation mechanism that triggers monthly. In addition, the system does not allow the journals to get larger than 10% of the file system that they are on, or leaving less than 15% of the file system free. You can modify these values for both the runtime and persistent journals in the `/etc/systemd/journald.conf` configuration file.

The `systemd-journald` process logs the current limits on the size of the journal when it starts. The following command output shows the journal entries that reflect the current size limits:

```
[user@host ~]$ journalctl | grep -E 'Runtime Journal|System Journal'
Mar 15 04:21:14 localhost systemd-journald[226]: Runtime Journal (/run/log/
journal/4ec03abd2f7b40118b1b357f479b3112) is 8.0M, max 113.3M, 105.3M free.
Mar 15 04:21:19 host.lab.example.com systemd-journald[719]: Runtime Journal (/run/
log/journal/4ec03abd2f7b40118b1b357f479b3112) is 8.0M, max 113.3M, 105.3M free.
Mar 15 04:21:19 host.lab.example.com systemd-journald[719]: System Journal (/run/
log/journal/4ec03abd2f7b40118b1b357f479b3112) is 8.0M, max 4.0G, 4.0G free.
```

**Note**

In the previous `grep` command, the vertical bar (`|`) symbol acts as an *or* operator. That is, the `grep` command matches any line with either the `Runtime Journal` string or the `System Journal` string from the `journalctl` command output. This command fetches the current size limits on the volatile (`Runtime`) journal store and on the persistent (`System`) journal store.

Configure Persistent System Journals

Configure the `systemd-journald` service as follows to preserve system journals persistently across a reboot:

- Create the `/var/log/journal` directory.

```
[root@host ~]# mkdir /var/log/journal
```

- Set the `Storage` parameter to the `persistent` value in the `/etc/systemd/journald.conf` file. Run your chosen text editor as the superuser to edit the `/etc/systemd/journald.conf` file.

```
[Journal]
Storage=persistent
...output omitted...
```

- Restart the `systemd-journald` service to apply the configuration changes.

```
[root@host ~]# systemctl restart systemd-journald
```

If the `systemd-journald` service successfully restarts, then the service creates subdirectories in the `/var/log/journal` directory. The subdirectory in the `/var/log/journal` directory has hexadecimal characters in its long name and contain files with the `.journal` extension. The `.journal` binary files store the structured and indexed journal entries.

```
[root@host ~]# ls /var/log/journal
4ec03abd2f7b40118b1b357f479b3112
[root@host ~]# ls /var/log/journal/4ec03abd2f7b40118b1b357f479b3112
system.journal  user-1000.journal
```

Although the system journals persist after a reboot, the `journalctl` command output includes entries from the current system boot as well as from the previous system boots. To limit the output to a specific system boot, use the `journalctl` command `-b` option. The following `journalctl` command retrieves the entries from the first system boot only:

```
[root@host ~]# journalctl -b 1
...output omitted...
```

The following `journalctl` command retrieves the entries from the second system boot only. The argument is meaningful only if the system was rebooted at least twice:

```
[root@host ~]# journalctl -b 2
...output omitted...
```

You can list the system boot events that the `journalctl` command recognizes, by using the `--list-boots` option.

```
[root@host ~]# journalctl --list-boots
-6 27de... Wed 2022-04-13 20:04:32 EDT-Wed 2022-04-13 21:09:36 EDT
-5 6a18... Tue 2022-04-26 08:32:22 EDT-Thu 2022-04-28 16:02:33 EDT
-4 e2d7... Thu 2022-04-28 16:02:46 EDT-Fri 2022-05-06 20:59:29 EDT
-3 45c3... Sat 2022-05-07 11:19:47 EDT-Sat 2022-05-07 11:53:32 EDT
-2 dfae... Sat 2022-05-07 13:11:13 EDT-Sat 2022-05-07 13:27:26 EDT
-1 e754... Sat 2022-05-07 13:58:08 EDT-Sat 2022-05-07 14:10:53 EDT
0 ee2c... Mon 2022-05-09 09:56:45 EDT-Mon 2022-05-09 12:57:21 EDT
```

The following `journalctl` command retrieves the entries from the current system boot only:

```
[root@host ~]# journalctl -b
...output omitted...
```

Note

When debugging a system crash with a persistent journal, usually you must limit the journal query to the reboot before the crash happened. You can use the `journalctl` command `-b` option with a negative number to indicate how many earlier system boots to include in the output. For example, the `journalctl -b -1` command limits the output to only the previous boot.

References

`systemd-journald.conf(5)`, `systemd-journald(8)` man pages

For more information, refer to the *Troubleshooting Problems Using Log Files* section in the *Red Hat Enterprise Linux 9 Configuring Basic System Settings* guide at https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/configuring_basic_system_settings/index#troubleshooting-problems-using-log-files_getting-started-with-system-administration

► Guided Exercise

Preserve the System Journal

Configure the system journal to preserve its data after a reboot.

Outcomes

- Configure the system journal to preserve its data after a reboot.

Before You Begin

As the **student** user on the **workstation** machine, use the **lab** command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start logs-preserve
```

Instructions

- 1. From the **workstation** machine, log in to the **servera** machine as the **student** user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
```

- 2. As the superuser, confirm that the **/var/log/journal** directory does not exist. Use the **ls** command to list the **/var/log/journal** directory contents. Use the **sudo** command to elevate the **student** user privileges. If prompted, use the **student** password.

```
[student@servera ~]$ sudo ls /var/log/journal
[sudo] password for student: student
ls: cannot access '/var/log/journal': No such file or directory
```

Because the **/var/log/journal** directory does not exist, the **systemd-journald** service does not preserve the log data after a reboot.

- 3. Configure the **systemd-journald** service on the **servera** machine to preserve journals after a reboot.

- 3.1. Create the **/var/log/journal** directory.

```
[student@servera ~]$ sudo mkdir /var/log/journal
```

- 3.2. Uncomment the **Storage=auto** line in the **/etc/systemd/journald.conf** file, and set the **Storage** parameter to the **persistent** value. You might use the **sudo vim /etc/systemd/journald.conf** command to edit the configuration file.

You can type `/Storage=auto` in the vim editor command mode to search for the `Storage=auto` line.

```
...output omitted...
[Journal]
Storage=persistent
...output omitted...
```

3.3. Restart the `systemd-journald` service to apply the configuration changes.

```
[student@servera ~]$ sudo systemctl restart systemd-journald.service
```

- ▶ 4. Verify that the `systemd-journald` service on the `servera` machine preserves its journals so that they persist after a reboot.

4.1. Restart the `servera` machine.

```
[student@servera ~]$ sudo systemctl reboot
Connection to servera closed by remote host.
Connection to servera closed.
[student@workstation ~]$
```

The SSH connection terminates as soon as you restart the `servera` machine.

4.2. Log in to the `servera` machine.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
```

- 4.3. Verify that a subdirectory with a long hexadecimal name exists in the `/var/log/journal` directory. You can find the journal files in that directory. The subdirectory name on your system might be different.

```
[student@servera ~]$ sudo ls /var/log/journal
[sudo] password for student: student
63b272eae8d5443ca7aaa5593479b25f
[student@servera ~]$ sudo ls /var/log/journal/63b272eae8d5443ca7aaa5593479b25f
system.journal  user-1000.journal
```

4.4. Return to the `workstation` system as the `student` user.

```
[student@servera ~]$ exit
logout
Connection to servera closed.
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish logs-preserve
```

Maintain Accurate Time

Objectives

- Maintain accurate time synchronization with Network Time Protocol (NTP) and configure the time zone to ensure correct time stamps for events recorded by the system journal and logs.

Administer Local Clocks and Time Zones

System time synchronization is critical for log file analysis across multiple systems. Also, some services might require time synchronization to work correctly. Machines use the Network Time Protocol to provide and obtain correct time information over the internet. A machine might get accurate time information from public NTP services, such as the NTP Pool Project. Another option is to sync with a high-quality hardware clock to serve accurate time to local clients.

The `timedatectl` command shows an overview of the current time-related system settings, including the current time, time zone, and NTP synchronization settings of the system.

```
[user@host ~]$ timedatectl
      Local time: Wed 2022-03-16 05:53:05 EDT
      Universal time: Wed 2022-03-16 09:53:05 UTC
            RTC time: Wed 2022-03-16 09:53:05
          Time zone: America/New_York (EDT, -0400)
System clock synchronized: yes
              NTP service: active
          RTC in local TZ: no
```

You can list a database of time zones with the `timedatectl` command `list-timezones` option:

```
[user@host ~]$ timedatectl list-timezones
Africa/Abidjan
Africa/Accra
Africa/Addis_Ababa
Africa/Algiers
Africa/Asmara
Africa/Bamako
...output omitted...
```

The Internet Assigned Numbers Authority (IANA) provides a public time zone database, and the `timedatectl` command bases the time zone names on that database. IANA names time zones based on the continent or ocean, and then typically (not always) the largest city within the time zone region. For example, most of the US Mountain time zone is `America/Denver`.

Some localities inside the time zone have different daylight saving time rules. For example, in the US, much of the state of Arizona (US Mountain time) does not change to daylight saving time, and is in the `America/Phoenix` time zone.

Use the `tzselect` command to identify the correct time zone name. This command interactively prompts the user with questions about the system's location, and outputs the name of the correct time zone. It does not change the system's time zone setting.

The `root` user can change the system setting to update the current time zone with the `timedatectl` command `set-timezone` option. For example, the following `timedatectl` command updates the current time zone to `America/Phoenix`.

```
[root@host ~]# timedatectl set-timezone America/Phoenix
[root@host ~]# timedatectl
          Local time: Wed 2022-03-16 03:05:55 MST
          Universal time: Wed 2022-03-16 10:05:55 UTC
             RTC time: Wed 2022-03-16 10:05:55
            Time zone: America/Phoenix (MST, -0700)
System clock synchronized: yes
              NTP service: active
          RTC in local TZ: no
```

**Note**

You can set a server's time zone to Coordinated Universal Time (UTC). The `tzselect` command does not include the name of the UTC time zone. Use the `timedatectl set-timezone UTC` command to set the system's current time zone to UTC.

Use the `timedatectl` command `set-time` option to change the system's current time. You might specify the time in the `"YYYY-MM-DD hh:mm:ss"` format, where you can omit either the date or the time. For example, the following `timedatectl` command changes the time to `09:00:00`.

```
[root@host ~]# timedatectl set-time 9:00:00
[root@host ~]# timedatectl
          Local time: Fri 2019-04-05 09:00:27 MST
          Universal time: Fri 2019-04-05 16:00:27 UTC
             RTC time: Fri 2019-04-05 16:00:27
            Time zone: America/Phoenix (MST, -0700)
System clock synchronized: yes
              NTP service: active
          RTC in local TZ: no
```

**Note**

The previous example might fail with the "Failed to set time: Automatic time synchronization is enabled" error message. In that case, first disable the automatic time synchronization before manually setting the date or time, as explained after this note.

The `timedatectl` command `set-ntp` option enables or disables NTP synchronization for automatic time adjustment. The option requires either a `true` or a `false` argument to turn it on or off. For example, the following `timedatectl` command turns off NTP synchronization.


```
[root@host ~]# timedatectl set-ntp false
```

**Note**

In Red Hat Enterprise Linux 9, the `timedatectl set-ntp` command adjusts whether the `chronyd` NTP service is enabled. Other Linux distributions might use this setting to adjust a different NTP or a *Simple Network Time Protocol (SNTP)* service.

Enabling or disabling NTP with other utilities in Red Hat Enterprise Linux, such as in the graphical GNOME Settings application, also updates this setting.

Configure and Monitor the chronyd Service

The `chronyd` service tracks the usually inaccurate local *Real-Time Clock (RTC)* by synchronizing it to the configured NTP servers. If no network connectivity is available, then the `chronyd` service calculates the RTC clock drift, and records it in the file that the `driftfile` value specifies in the `/etc/chrony.conf` configuration file.

By default, the `chronyd` service uses servers from the NTP Pool Project to synchronize time and requires no additional configuration. You might need to change the NTP servers for a machine that runs on an isolated network.

The *stratum* of the NTP time source determines its quality. The stratum determines the number of hops that the machine is away from a high-performance reference clock. The reference clock is a `stratum 0` time source. An NTP server that is directly attached to the reference clock is a `stratum 1` time source. A machine that synchronizes time from the NTP server is a `stratum 2` time source.

The `server` and the `peer` are the two categories of time sources that you can declare in the `/etc/chrony.conf` configuration file. The `server` is one stratum above the local NTP server, and the `peer` is at the same stratum level. You can define multiple servers and peers in the `chronyd` configuration file, one per line.

The first argument of the `server` line is the IP address or DNS name of the NTP server. Following the server IP address or name, you can list a series of options for the server. Red Hat recommends using the `iburst` option, because then the `chronyd` service takes four measurements in a short time period for a more accurate initial clock synchronization after the service starts. For more information about the `chronyd` configuration file options, use the `man 5 chrony.conf` command.

As an example, with the following `server classroom.example.com iburst` line in the `/etc/chrony.conf` configuration file, the `chronyd` service uses the `classroom.example.com` server as the NTP time source.

```
# Use public servers from the pool.ntp.org project.
...output omitted...
server classroom.example.com iburst
...output omitted...
```

Restart the service after pointing the `chronyd` service to the `classroom.example.com` local time source.

```
[root@host ~]# systemctl restart chronyd
```

The `chronyc` command acts as a client to the `chronyd` service. After setting up NTP synchronization, verify that the local system is seamlessly using the NTP server to synchronize the system clock, by using the `chronyc sources` command. For more verbose output with additional explanations about the output, use the `chronyc sources -v` command.

```
[root@host ~]# chronyc sources -v

.-- Source mode  '^' = server, '=' = peer, '#' = local clock.
/ .- Source state '*' = current best, '+' = combined, '-' = not combined,
| /           'x' = may be in error, '~' = too variable, '?' = unusable.
||
||                                     .- xxxx [ yyyy ] +/- zzzz
||      Reachability register (octal) -.      | xxxx = adjusted offset,
||      Log2(Polling interval) --.      |      | yyyy = measured offset,
||                                     \      |      | zzzz = estimated error.
||                                     |      |      \
MS Name/IP address             Stratum Poll Reach LastRx Last sample
=====
^* 172.25.254.254                3      6   17   26  +2957ns[+2244ns] +/-   25ms
```

The asterisk character (*) in the S (Source state) field indicates that the `chronyd` service uses the `classroom.example.com` server as a time source and is the NTP server that the machine is currently synchronized to.



References

`timedatectl(1)`, `tzselect(8)`, `chronyd(8)`, `chrony.conf(5)`, and `chronyc(1)` man pages

NTP Pool Project

<http://www.ntppool.org/>

Time Zone Database

<http://www.iana.org/time-zones>

► Guided Exercise

Maintain Accurate Time

Adjust the time zone on a server and ensure that its system clock is synchronized with an NTP time source.

Outcomes

- Change the time zone on a server.
- Configure the server to synchronize its time with an NTP time source.

Before You Begin

As the `student` user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start logs-maintain
```

Instructions

- 1. Log in to the `servera` machine as the `student` user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
```

- 2. For this exercise, pretend that the `servera` machine is relocated to Haiti and that you need to update the time zone. Elevate the privileges of the `student` user to run the `timedatectl` command to update the time zone.

2.1. Select the appropriate time zone for Haiti.

```
[root@servera ~]# 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 timezone using the Posix TZ format.
#? 2
Please select a country whose clocks agree with yours.
 1) Anguilla      19) Dominican Republic  37) Peru
 2) Antigua & Barbuda  20) Ecuador      38) Puerto Rico
 3) Argentina    21) El Salvador    39) St Barthelemy
 4) Aruba        22) French Guiana  40) St Kitts & Nevis
 5) Bahamas     23) Greenland     41) St Lucia
 6) Barbados    24) Grenada       42) St Maarten (Dutch)
 7) Belize      25) Guadeloupe    43) St Martin (French)
 8) Bolivia     26) Guatemala     44) St Pierre & Miquelon
 9) Brazil      27) Guyana        45) St Vincent
10) Canada      28) Haiti         46) Suriname
11) Caribbean NL  29) Honduras      47) Trinidad & Tobago
12) Cayman Islands  30) Jamaica      48) Turks & Caicos Is
13) Chile       31) Martinique   49) United States
14) Colombia    32) Mexico        50) Uruguay
15) Costa Rica  33) Montserrat    51) Venezuela
16) Cuba        34) Nicaragua     52) Virgin Islands (UK)
17) Curaçao     35) Panama        53) Virgin Islands (US)
18) Dominica    36) Paraguay
#? 28

```

The following information has been given:

Haiti

Therefore TZ='America/Port-au-Prince' will be used.
 Selected time is now: Wed Mar 16 07:10:35 EDT 2022.
 Universal Time is now: Wed Mar 16 11:10:35 UTC 2022.
 Is the above information OK?

```

1) Yes
2) No
#? 1

```

You can make this change permanent for yourself by appending the line
 TZ='America/Port-au-Prince'; 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:
 America/Port-au-Prince

2.2. Update the time zone on the servera machine to America/Port-au-Prince.

```

[root@servera ~]# timedatectl set-timezone \
  America/Port-au-Prince

```

2.3. Verify that you correctly set the time zone to America/Port-au-Prince.

```
[root@servera ~]# timedatectl
      Local time: Wed 2022-03-16 07:13:25 EDT
      Universal time: Wed 2022-03-16 11:13:25 UTC
          RTC time: Wed 2022-03-16 11:13:24
      Time zone: America/Port-au-Prince (EDT, -0400)
System clock synchronized: no
      NTP service: inactive
      RTC in local TZ: no
```

- ▶ **3.** Configure the `chronyd` service on the `servera` machine to synchronize the system time with the `classroom.example.com` server as the NTP time source.

- 3.1. Edit the `/etc/chrony.conf` configuration file to specify the `classroom.example.com` server as the NTP time source. The following output shows the configuration line to add to the configuration file, which includes the `iburst` option to speed up initial time synchronization:

```
...output omitted...
server classroom.example.com iburst
...output omitted...
```

- 3.2. Enable time synchronization on the `servera` machine. The command activates the NTP server with the settings from the `/etc/chrony.conf` configuration file. That command might activate either the `chronyd` or the `ntpd` service, depending on which service is currently installed on the system.

```
[root@servera ~]# timedatectl set-ntp true
```

- ▶ **4.** Verify that the `servera` machine configuration synchronizes with the `classroom.example.com` time source in the classroom environment.
- 4.1. Verify that time synchronization is enabled on the `servera` machine.

**Note**

If the output shows that the clock is not synchronized, then wait for a few seconds and rerun the `timedatectl` command. It takes a few seconds to successfully synchronize the time settings with the time source.

```
[root@servera ~]# timedatectl
      Local time: Wed 2022-03-16 07:24:13 EDT
      Universal time: Wed 2022-03-16 11:24:13 UTC
          RTC time: Wed 2022-03-16 11:24:13
      Time zone: America/Port-au-Prince (EDT, -0400)
System clock synchronized: yes
      NTP service: active
      RTC in local TZ: no
```

- 4.2. Verify that the `servera` machine currently synchronizes its time settings with the `classroom.example.com` time source.

The output shows an asterisk character (*) in the source state (S) field for the `classroom.example.com` NTP time source. The asterisk indicates that the local system time successfully synchronizes with the NTP time source.

```
[root@servera ~]# chronyc sources -v

.-- Source mode  '^' = server, '=' = peer, '#' = local clock.
/ .- Source state '*' = current best, '+' = combined, '-' = not combined,
| /           'x' = may be in error, '~' = too variable, '?' = unusable.
||
||                                     .- xxxx [ yyyy ] +/- zzzz
||      Reachability register (octal) --.    | xxxx = adjusted offset,
||      Log2(Polling interval) --.    |    | yyyy = measured offset,
||                                     \    |    | zzzz = estimated error.
||                                     |    | \
||                                     |    |
MS Name/IP address             Stratum Poll Reach LastRx Last sample
=====
^* classroom.example.com         3      6   17    51  -118ns[ +49us] +/-  25ms
```

4.3. Return to the workstation system as the student user.

```
[root@servera ~]# exit
logout
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the `workstation` machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish logs-maintain
```

► Lab

Analyze and Store Logs

Change the time zone on an existing server and configure a new log file for all events for authentication failures.

Outcomes

- Update the time zone on an existing server.
- Configure a new log file to store all messages for authentication failures.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start logs-review
```

Instructions

1. Log in to the `serverb` machine as the `student` user.
2. Pretend that the `serverb` machine is relocated to Jamaica and that you must update the time zone to `America/Jamaica`. Verify that you correctly set the appropriate time zone.
3. View the recorded log events in the previous 30 minutes on the `serverb` machine.
4. Create the `/etc/rsyslog.d/auth-errors.conf` file. Configure the `rsyslog` service to write the `Logging test authpriv.alert` message to the `/var/log/auth-errors` file. Use the `authpriv` facility and the `alert` priority.

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade logs-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish logs-review
```

► Solution

Analyze and Store Logs

Change the time zone on an existing server and configure a new log file for all events for authentication failures.

Outcomes

- Update the time zone on an existing server.
- Configure a new log file to store all messages for authentication failures.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start logs-review
```

Instructions

1. Log in to the `serverb` machine as the `student` user.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
[student@serverb ~]$
```

2. Pretend that the `serverb` machine is relocated to Jamaica and that you must update the time zone to `America/Jamaica`. Verify that you correctly set the appropriate time zone.

- 2.1. Select the appropriate time zone for Jamaica.

```
[student@serverb ~]$ 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 timezone using the Posix TZ format.
#? 2
```


Please select a country whose clocks agree with yours.

1) Anguilla	19) Dominican Republic	37) Peru
2) Antigua & Barbuda	20) Ecuador	38) Puerto Rico
3) Argentina	21) El Salvador	39) St Barthelemy
4) Aruba	22) French Guiana	40) St Kitts & Nevis
5) Bahamas	23) Greenland	41) St Lucia
6) Barbados	24) Grenada	42) St Maarten (Dutch)
7) Belize	25) Guadeloupe	43) St Martin (French)
8) Bolivia	26) Guatemala	44) St Pierre & Miquelon
9) Brazil	27) Guyana	45) St Vincent
10) Canada	28) Haiti	46) Suriname
11) Caribbean NL	29) Honduras	47) Trinidad & Tobago
12) Cayman Islands	30) Jamaica	48) Turks & Caicos Is
13) Chile	31) Martinique	49) United States
14) Colombia	32) Mexico	50) Uruguay
15) Costa Rica	33) Montserrat	51) Venezuela
16) Cuba	34) Nicaragua	52) Virgin Islands (UK)
17) Curaçao	35) Panama	53) Virgin Islands (US)
18) Dominica	36) Paraguay	

#? **30**

The following information has been given:

Jamaica

Therefore TZ='America/Jamaica' will be used.

Selected time is now: Wed Mar 16 07:17:15 EST 2022.

Universal Time is now: Wed Mar 16 12:17:15 UTC 2022.

Is the above information OK?

1) Yes

2) No

#? **1**

You can make this change permanent for yourself by appending the line

TZ='America/Jamaica'; 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:

America/Jamaica

- 2.2. Elevate the student user privileges to update the time zone of the serverb server to America/Jamaica.

```
[student@serverb ~]$ sudo timedatectl set-timezone America/Jamaica
[sudo] password for student: student
```

- 2.3. Verify that you successfully set the time zone to America/Jamaica.

```
[student@serverb ~]$ timedatectl
      Local time: Wed 2022-03-16 07:18:40 EST
      Universal time: Wed 2022-03-16 12:18:40 UTC
      RTC time: Wed 2022-03-16 12:18:40
      Time zone: America/Jamaica (EST, -0500)
System clock synchronized: yes
      NTP service: active
      RTC in local TZ: no
```

3. View the recorded log events in the previous 30 minutes on the `serverb` machine.

```
[student@serverb ~]$ journalctl --since "-30min"
...output omitted...
Mar 16 07:10:58 localhost kernel: x86/PAT: Configuration [0-7]: WB WC UC- UC WB
WP UC- WT
Mar 16 07:10:58 localhost kernel: found SMP MP-table at [mem
0x000f5bd0-0x000f5bdf]
Mar 16 07:10:58 localhost kernel: Using GB pages for direct mapping
Mar 16 07:10:58 localhost kernel: RAMDISK: [mem 0x2e0d9000-0x33064fff]
Mar 16 07:10:58 localhost kernel: ACPI: Early table checksum verification disabled
Mar 16 07:10:58 localhost kernel: ACPI: RSDP 0x000000000000F5B9 000014 (v00
BOCHS )
Mar 16 07:10:58 localhost kernel: ACPI: RSDT 0x000000007FFE12C4 00002C (v01 BOCHS
BXPCRSDT 00000001 BXPC 00000001)
Mar 16 07:10:58 localhost kernel: ACPI: FACP 0x000000007FFE11D0 000074 (v01 BOCHS
BXPCFACP 00000001 BXPC 00000001)
Mar 16 07:10:58 localhost kernel: ACPI: DSDT 0x000000007FFDFDC0 001410 (v01 BOCHS
BXPCDSDT 00000001 BXPC 00000001)
lines 1-50/50 q
[student@serverb ~]$
```

4. Create the `/etc/rsyslog.d/auth-errors.conf` file. Configure the `rsyslog` service to write the `Logging test authpriv.alert` message to the `/var/log/auth-errors` file. Use the `authpriv` facility and the `alert` priority.
 - 4.1. Create the `/etc/rsyslog.d/auth-errors.conf` file and specify the new `/var/log/auth-errors` file as the destination for authentication and security messages.

```
authpriv.alert /var/log/auth-errors
```

- 4.2. Restart the `rsyslog` service to apply the configuration file changes.

```
[student@serverb ~]$ sudo systemctl restart rsyslog
```

- 4.3. Use the `logger -p` command to write the `Logging test authpriv.alert` message to the `/var/log/auth-errors` file. Use the `authpriv` facility and the `alert` priority.

```
[student@serverb ~]$ logger -p authpriv.alert "Logging test authpriv.alert"
```

- 4.4. Verify that the `/var/log/auth-errors` file contains the log entry with the `Logging test authpriv.alert` message.

```
[student@serverb ~]$ sudo tail /var/log/auth-errors  
Mar 16 07:25:12 serverb student[1339]: Logging test authpriv.alert
```

4.5. Return to the workstation system as the student user.

```
[student@serverb ~]$ exit  
logout  
Connection to serverb closed.  
[student@workstation ~]$
```

Evaluation

As the student user on the workstation machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade logs-review
```

Finish

On the workstation machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish logs-review
```

Summary

- The `systemd-journald` and `rsyslog` services capture and write log messages to the appropriate files.
- The `/var/log` directory contains log files.
- Periodic rotation of log files prevents them from filling up the file-system space.
- The `systemd` journals are temporary and do not persist across a reboot.
- The `chronyd` service helps to synchronize time settings with a time source.
- You can update the time zone of the server based on its location.

Chapter 4

Archive and Transfer Files

Goal

Archive and copy files from one system to another.

Objectives

- Archive files and directories into a compressed file with `tar`, and extract the contents of an existing `tar` archive.
- Transfer files to or from a remote system securely with `SSH`.
- Efficiently and securely synchronize the contents of a local file or directory with a remote server copy.

Sections

- Manage Compressed `tar` Archives (and Guided Exercise)
- Transfer Files Between Systems Securely (and Guided Exercise)
- Synchronize Files Between Systems Securely (and Guided Exercise)

Lab

- Archive and Transfer Files

Manage Compressed tar Archives

Objectives

- Archive files and directories into a compressed file with `tar`, and extract the contents of an existing `tar` archive.

Create Archives from the Command Line

An *archive* is a single regular file or device file that contains multiple files. The device file could be a tape drive, flash drive, or other removable media. When using a regular file, archiving is analogous to the `zip` utility and similar variations that are popular on most operating systems.



Note

The original, ubiquitous `zip` compression and file packaging utility uses the PKZIP (Phil Katz's ZIP for MSDOS systems) algorithm, which has evolved significantly, and is supported on RHEL with the `zip` and `unzip` commands. Many other compression algorithms have been developed since `zip` was introduced, and each has its advantages. For creating compressed archives for general use, any `tar`-supported compression algorithm is acceptable.

Archive files are used to create manageable personal backups, or to simplify transferring a set of files across a network when other methods, such as `rsync`, are unavailable or might be more complex. Archive files can be created with or without using compression to reduce the archive file size.

On Linux, the `tar` utility is the common command to create, manage, and extract archives. Use the `tar` command to gather multiple files into a single archive file. A *tar archive* is a structured sequence of file metadata and data with an index so you can extract individual files.

Files can be compressed during creation by using one of the supported compression algorithms. The `tar` command can list the contents of an archive without extracting, and can extract original files directly from both compressed and uncompressed archives.

Options of the tar Utility

One of the following `tar` command actions is required to perform a `tar` operation:

- **-c** or **--create** : Create an archive file.
- **-t** or **--list** : List the contents of an archive.
- **-x** or **--extract** : Extract an archive.

The following `tar` command general options are often included:

- **-v** or **--verbose** : Show the files that are being archived or extracted during the `tar` operation.
- **-f** or **--file** : Follow this option with the archive file name to create or open.
- **-p** or **--preserve-permissions** : Preserve the original file permissions when extracting.
- **--xattrs** : Enable extended attribute support, and store extended file attributes.

- **--selinux** : Enable SELinux context support, and store SELinux file contexts.

The following **tar** command compression options are used to select an algorithm:

- **-a** or **--auto-compress** : Use the archive's suffix to determine the algorithm to use.
- **-z** or **--gzip** : Use the gzip compression algorithm, which results in a **.tar.gz** suffix.
- **-j** or **--bzip2** : Use the bzip2 compression algorithm, which results in a **.tar.bz2** suffix.
- **-J** or **--xz** : Use the xz compression algorithm, which results in a **.tar.xz** suffix.



Note

The **tar** command still supports the legacy option style that does not use a dash (-) character. You might find this syntax in legacy scripts or documentation, and the behavior is essentially the same. For command consistency, Red Hat recommends using the short- or long-option styles instead.

Create an Archive

To create an archive with the **tar** command, use the **create** and **file** options with the archive file name as the first argument, followed by a list of files and directories to include in the archive.

The **tar** command recognizes absolute and relative file name syntax. By default, **tar** removes the leading forward slash (/) character from absolute file names, so that files are stored internally with relative path names. This technique is safer, because extracting absolute path names always overwrites existing files. With files that are archived with relative path names, files can be extracted to a new directory without overwriting existing files.

The following command creates the **mybackup.tar** archive to contain the **myapp1.log**, **myapp2.log**, and **myapp3.log** files from the user's home directory. If a file with the same name as the requested archive exists in the target directory, then the **tar** command overwrites the file.

```
[user@host ~]$ tar -cf mybackup.tar myapp1.log myapp2.log myapp3.log
[user@host ~]$ ls mybackup.tar
mybackup.tar
```

A user must have read permissions on the target files that are being archived. For example, creating an archive in the **/etc** directory requires **root** privileges, because only privileged users can read all **/etc** files. An unprivileged user can create an archive of the **/etc** directory, but the archive excludes files that the user cannot read, and directories for which the user lacks the read and execute permissions.

In this example, the **root** user creates the **/root/etc-backup.tar** archive of the **/etc** directory.

```
[root@host ~]# tar -cf /root/etc-backup.tar /etc
tar: Removing leading '/' from member names
```



Important

Extended file attributes, such as access control lists (ACL) and SELinux file contexts, are not preserved by default in an archive. Use the **--acls**, **--selinux**, and **--xattrs** options to include POSIX ACLs, SELinux file contexts, and other extended attributes, respectively.

List Archive Contents

Use the `tar` command `t` option to list the file names from within the archive that are specified with the `f` option. The files list with relative name syntax, because the leading forward slash was removed during archive creation.

```
[root@host ~]# tar -tf /root/etc.tar
etc/
etc/fstab
etc/crypttab
etc/mtab
...output omitted...
```

Extract Archive Contents

Extract a `tar` archive into an empty directory to avoid overwriting existing files. When the `root` user extracts an archive, the extracted files preserve the original user and group ownership. If a regular user extracts files, then the user becomes the owner of the extracted files.

List the contents of the `/root/etc.tar` archive and then extract its files to the `/root/etcbakup` directory:

```
[root@host ~]# mkdir /root/etcbakup
[root@host ~]# cd /root/etcbakup
[root@host etcbakup]# tar -tf /root/etc.tar
etc/
etc/fstab
etc/crypttab
etc/mtab
...output omitted...
[root@host etcbakup]# tar -xf /root/etc.tar
```

When you extract files from an archive, the current `umask` is used to modify each extracted file's permissions. Instead, use the `tar` command `p` option to preserve the original archived permissions for extracted files. The `--preserve-permissions` option is enabled by default for a superuser.

```
[user@host scripts]# tar -xpf /home/user/myscripts.tar
...output omitted...
```

Create a Compressed Archive

The `tar` command supports these compression methods, and others:

- **gzip** compression is the earlier, fastest method, and is widely available across platforms.
- **bzip2** compression creates smaller archives but is less widely available than `gzip`.
- **xz** compression is newer, and offers the best compression ratio of the available methods.

The effectiveness of any compression algorithm depends on the type of data that is compressed. Previously compressed data files, such as picture formats or RPM files, typically do not significantly compress further.

Create the `/root/etcbakup.tar.gz` archive with `gzip` compression from the contents of the `/etc` directory:

```
[root@host ~]# tar -czf /root/etcbbackup.tar.gz /etc
tar: Removing leading `/' from member names
```

Create the `/root/logbackup.tar.bz2` archive with `bzip2` compression from the contents of the `/var/log` directory:

```
[root@host ~]$ tar -cjf /root/logbackup.tar.bz2 /var/log
tar: Removing leading `/' from member names
```

Create the `/root/sshconfig.tar.xz` archive with `xz` compression from the contents of the `/etc/ssh` directory:

```
[root@host ~]$ tar -cjf /root/sshconfig.tar.xz /etc/ssh
tar: Removing leading `/' from member names
```

After creating an archive, verify its table of contents with the `tar` command `tf` options. It is not necessary to specify the compression option when listing a compressed archive file, because the compression type is read from the archive's header. List the archived content in the `/root/etcbbackup.tar.gz` file, which uses the `gzip` compression:

```
[root@host ~]# tar -tf /root/etcbbackup.tar.gz
etc/
etc/fstab
etc/crypttab
etc/mtab
...output omitted...
```

Extract Compressed Archive Contents

The `tar` command can automatically determine which compression was used, so it is not necessary to specify the compression option. If you do include an incorrect compression type, `tar` reports that the specified compression type does not match the file's type. In the following example, the `tar` command uses the `-z` option, which indicates `gzip` compression, but the file name extension is `.xz`, which indicates `xz` compression:

```
[root@host ~]# tar -xzf /root/etcbbackup.tar.xz

gzip: stdin: not in gzip format
tar: Child returned status 1
tar: Error is not recoverable: exiting now
```

Listing a compressed `tar` archive works in the same way as listing an uncompressed `tar` archive. Use the `tar` command with the `tf` option to verify the content of the compressed archive before extracting its contents:

```
[root@host logbackup]# tar -tf /root/logbackup.tar
var/log/
var/log/lastlog
var/log/README
var/log/private/
...output omitted...
```

The `gzip`, `bzip2`, and `xz` algorithms are also implemented as stand-alone commands for compressing individual files without creating an archive. With these commands, you cannot create a single compressed file of multiple files, such as a directory. As previously discussed, to create a compressed archive of multiple files, use the `tar` command with your preferred compression option. To uncompress a single compressed file or a compressed archive file without extracting its contents, use the `gunzip`, `bunzip2`, and `unxz` stand-alone commands.

The `gzip` and `xz` commands provide an `-l` option to view the uncompressed size of a compressed single or archive file. Use this option to verify that enough space is available before uncompressing or extracting a file.

```
[user@host ~]$ gzip -l file.tar.gz
      compressed      uncompressed   ratio uncompressed_name
      221603125        303841280   27.1% file.tar

[user@host ~]$ xz -l file.xz
Strms  Blocks   Compressed Uncompressed  Ratio  Check  Filename
    1      1    195.7 MiB   289.8 MiB   0.675  CRC64   file.xz
```



References

`tar(1)`, `gzip(1)`, `gunzip(1)`, `bzip2(1)`, `bunzip2(1)`, `xz(1)`, and `unxz(1)` man pages

► Guided Exercise

Manage Compressed tar Archives

Create archive files and extract their contents with the `tar` command.

Outcomes

- Archive a directory tree and extract the archive content to another location.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command ensures that all required resources are available.

```
[student@workstation ~]$ lab start archive-manage
```

Instructions

- 1. From `workstation`, log in to `servera` as the `student` user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

- 2. Create an archive of the `/etc` directory with `gzip` compression. Save the archive file as `/tmp/etc.tar.gz`.

```
[root@servera ~]# tar -czf /tmp/etc.tar.gz /etc
tar: Removing leading `/' from member names
```

- 3. Verify that the `etc.tar.gz` archive contains the files from the `/etc` directory.

```
[root@servera ~]# tar -tzf /tmp/etc.tar.gz
etc/
etc/mtab
etc/fstab
etc/crypttab
etc/resolv.conf
...output omitted...
```

- 4. Create the `/backuptest` directory. Verify that the `etc.tar.gz` backup file is a valid archive by decompressing the file to the `/backuptest` directory.

- 4.1. Create the `/backuptest` directory and change to that directory.

```
[root@servera ~]# mkdir /backuptest
[root@servera ~]# cd /backuptest
[root@servera backuptest]#
```

4.2. Extract the `/tmp/etc.tar.gz` archive to the `/backuptest` directory.

```
[root@servera backuptest]# tar -xzf /tmp/etc.tar.gz
```

4.3. List the contents of the `/backuptest` directory. Verify that the directory contains the `/etc` directory backup files.

```
[root@servera backuptest]# ls -l
total 12
drwxr-xr-x. 95 root root 8192 Feb  8 10:16 etc
[root@servera backuptest]# ls -l etc
total 1228
-rw-r--r--.  1 root root      12 Feb 24 05:25 adjtime
-rw-r--r--.  1 root root    1529 Jun 23  2020 aliases
drwxr-xr-x.  2 root root   4096 Mar  3 04:48 alternatives
...output omitted...
```

► 5. Return to the workstation system as the student user.

```
[root@servera backuptest]# exit
logout
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation]$
```

Finish

On the `workstation` machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish archive-manage
```

Transfer Files Between Systems Securely

Objectives

- Transfer files to or from a remote system securely with SSH.

Transfer Remote Files with the Secure File Transfer Program

The **OpenSSH** suite securely runs shell commands on remote systems. Use the *Secure File Transfer Program (SFTP)* to interactively upload to or download files from an SSH server. This program is part of the **OpenSSH** suite. A session with the **sftp** command uses the secure authentication mechanism and encrypted data transfer to and from the SSH server.

Specify a remote location for the source or destination of the files to copy. For the format of the remote location, use `[user@]host:/path`. The `user@` part of the argument is optional. If this part is missing, then the **sftp** command uses your current local username. When you run the **sftp** command, your terminal provides an **sftp>** prompt.

```
[user@host ~]$ sftp remoteuser@remotehost
remoteuser@remotehost's password: password
Connected to remotehost.
sftp>
```

The interactive **sftp** session accepts various commands that work the same way in the remote file system as in the local file system, such as the **ls**, **cd**, **mkdir**, **rmdir**, and **pwd** commands. The **put** command uploads a file to the remote system. The **get** command downloads a file from the remote system. The **exit** command exits the **sftp** session.

List the available **sftp** commands by using the **help** command in the **sftp** session:

```
sftp> help
Available commands:
bye                               Quit sftp
cd path                           Change remote directory to 'path'
chgrp [-h] grp path              Change group of file 'path' to 'grp'
chmod [-h] mode path             Change permissions of file 'path' to 'mode'
chown [-h] own path              Change owner of file 'path' to 'own'
...output omitted...
```

In an **sftp** session, you might run some commands on your local host. For most available commands, add the **l** character before the command. For example, the **pwd** command prints the current working directory on the remote host. To print the current working directory on your local host, use the **lpwd** command.

```
sftp> pwd
Remote working directory: /home/remoteuser
sftp> lpwd
Local working directory: /home/user
```

The next example uploads the `/etc/hosts` file on the local system to the newly created `/home/remoteuser/hostbackup` directory on the `remotehost` machine. The `sftp` session expects that the `put` command is followed by a local file in the connecting user's home directory, in this case the `/home/remoteuser` directory:

```
sftp> mkdir hostbackup
sftp> cd hostbackup
sftp> put /etc/hosts
Uploading /etc/hosts to /home/remoteuser/hostbackup/hosts
/etc/hosts                                100% 227      0.2KB/s   00:00
```

To copy a whole directory tree recursively, use the `sftp` command `-r` option. The following example recursively copies the `/home/user/directory` local directory to the `remotehost` machine.

```
sftp> put -r directory
Uploading directory/ to /home/remoteuser/directory
Entering directory/
file1                                100%   0      0.0KB/s   00:00
file2                                100%   0      0.0KB/s   00:00
sftp> ls -l
drwxr-xr-x  2 student  student    32 Mar 21 07:51 directory
```

To download the `/etc/yum.conf` file from the remote host to the current directory on the local system, execute the `get /etc/yum.conf` command, and then exit the `sftp` session.

```
sftp> get /etc/yum.conf
Fetching /etc/yum.conf to yum.conf
/etc/yum.conf                          100% 813      0.8KB/s   00:00
sftp> exit
[user@host ~]$
```

To get a remote file with the `sftp` command on a single command line, without opening an interactive session, use the following syntax. You cannot use single command-line syntax to put files on a remote host.

```
[user@host ~]$ sftp remoteuser@remotehost:/home/remoteuser/remotefile
Connected to remotehost.
Fetching /home/remoteuser/remotefile to remotefile
remotefile                                100%   7
15.7KB/s   00:00
```

Transfer Files with Secure Copy Protocol



Warning

In versions prior to RHEL 9, the `scp` command was based on a historical `rcp` protocol that was not designed with security considerations. The `scp` protocol has a known code injection issue such that an attacker could execute arbitrary commands on the remote server. For this reason, the `scp` protocol is not covered in this course.

Although some vulnerabilities with the `scp` protocol were fixed in recent years, not all can be fixed while maintaining backward compatibility. For this reason, Red Hat recommends no longer using the legacy `scp` protocol in new applications or scripts.

You can find more information about this issue in <https://access.redhat.com/security/cve/cve-2020-15778>.

The `scp` Secure Copy command, which is also part of the `OpenSSH` suite, copies files from a remote system to the local system, or from the local system to a remote system. Since RHEL 9, the `scp` command uses the `sftp` protocol to transfer files.

You can specify a remote location for the source or destination of the files that you are copying. As with the `sftp` command, the `scp` command uses `[user@]host` to identify the target system and username. If you do not specify a user, then the command attempts to log in with your local username as the remote username.

To revert to the legacy `scp` protocol, you can include the `-O` flag. Due to the code injection vulnerability with the `scp` protocol, Red Hat recommends that you do not use the `-O` flag when you use the `scp` command.



References

`sftp(1)` man pages

► Guided Exercise

Transfer Files Between Systems Securely

Copy files from a remote system to a local directory with `sftp`.

Outcomes

- Copy files from a remote host to a directory on the local machine.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start archive-transfer
```

Instructions

- 1. Use the `ssh` command to log in to `servera` as the `student` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
```

- 2. Use the `sftp` command to copy the `/etc/ssh` directory from the `serverb` machine to the `/home/student/serverbackup` directory on the `servera` machine.

- 2.1. On the `servera` machine, create a `/home/student/serverbackup` directory.

```
[student@servera ~]$ mkdir ~/serverbackup
```

- 2.2. Use the `sftp` command to open a session to the `serverb` machine. Only the `root` user can read all the content in the `/etc/ssh` directory. When prompted, enter `redhat` as the password.

```
[student@servera ~]$ sftp root@serverb
root@serverb's password: redhat
Connected to serverb.
sftp>
```

- 2.3. Change the local current directory to the newly created `/home/student/serverbackup` directory.

```
sftp> lcd /home/student/serverbackup/
sftp> lpwd
Local working directory: /home/student/serverbackup
```

2.4. Recursively copy the /etc/ssh directory from the serverb machine to the /home/student/serverbackup directory on the servera machine.

```
sftp> get -r /etc/ssh
Fetching /etc/ssh/ to ssh
Retrieving /etc/ssh
Retrieving /etc/ssh/sshd_config.d
50-redhat.conf                               100% 719   881.5KB/s   00:00
Retrieving /etc/ssh/ssh_config.d
50-redhat.conf                               100% 581   347.4KB/s   00:00
01-training.conf                             100% 36    25.8KB/s    00:00
moduli                                       100% 565KB 71.9MB/s    00:00
ssh_config                                  100% 1921   1.1MB/s     00:00
ssh_host_rsa_key                            100% 2602   7.2MB/s     00:00
ssh_host_rsa_key.pub                        100% 565    1.6MB/s     00:00
ssh_host_ecdsa_key                          100% 505    1.6MB/s     00:00
ssh_host_ecdsa_key.pub                      100% 173   528.6KB/s   00:00
ssh_host_ed25519_key                        100% 399    1.0MB/s     00:00
ssh_host_ed25519_key.pub                    100% 93    275.8KB/s   00:00
sshd_config                                100% 3730   10.3MB/s    00:00
```

2.5. Exit from the sftp session. Verify that the /etc/ssh directory from the serverb machine is copied to the /home/student/serverbackup directory on the servera machine.

```
sftp> exit
[student@servera ~]$ ls -lR ~/serverbackup
/home/student/serverbackup:
total 4
drwxr-xr-x. 4 student student 4096 Mar 21 12:01 ssh

/home/student/serverbackup/ssh:
total 600
-rw-r--r--. 1 student student 578094 Mar 21 12:01 moduli
-rw-r--r--. 1 student student 1921 Mar 21 12:01 ssh_config
drwxr-xr-x. 2 student student 52 Mar 21 12:01 ssh_config.d
-rw-----. 1 student student 3730 Mar 21 12:01 sshd_config
drwx-----. 2 student student 28 Mar 21 12:01 sshd_config.d
-rw-r-----. 1 student student 505 Mar 21 12:01 ssh_host_ecdsa_key
-rw-r--r--. 1 student student 173 Mar 21 12:01 ssh_host_ecdsa_key.pub
-rw-r-----. 1 student student 399 Mar 21 12:01 ssh_host_ed25519_key
-rw-r--r--. 1 student student 93 Mar 21 12:01 ssh_host_ed25519_key.pub
-rw-r-----. 1 student student 2602 Mar 21 12:01 ssh_host_rsa_key
-rw-r--r--. 1 student student 565 Mar 21 12:01 ssh_host_rsa_key.pub

/home/student/serverbackup/ssh/ssh_config.d:
total 8
-rw-r--r--. 1 student student 36 Mar 21 12:01 01-training.conf
```

```
-rw-r--r--. 1 student student 581 Mar 21 12:01 50-redhat.conf  
  
/home/student/serverbackup/ssh/sshd_config.d:  
total 4  
-rw-----. 1 student student 719 Mar 21 12:01 50-redhat.conf
```

3. Return to the workstation system as the student user.

```
[student@servera ~]$ exit  
logout  
Connection to servera closed.  
[student@workstation]$
```

Finish

On the workstation machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish archive-transfer
```

Synchronize Files Between Systems Securely

Objectives

- Efficiently and securely synchronize the contents of a local file or directory with a remote server copy.

Synchronize Remote Files and Directories

The `rsync` command is another way to copy files from one system to another system securely. The tool uses an algorithm that minimizes copied data by synchronizing only the changed portions of files. If two files or directories are similar between two servers, then the `rsync` command copies only the differences between the file systems.

An advantage of the `rsync` command is that it copies files securely and efficiently between a local system and a remote system. Whereas an initial directory synchronization takes about the same time as copying it, subsequent synchronizations copy only the differences over the network, which substantially accelerates updates.

Use the `rsync` command `-n` option for a dry run. A dry run simulates what happens when the command is executed. The dry run shows the changes that the `rsync` command would perform when executing the command. Perform a dry run before the actual `rsync` command operation to ensure that no critical files are overwritten or deleted.

When synchronizing with the `rsync` command, two standard options are the `-v` and `-a` options.

The `rsync` command `-v` or `--verbose` option provides a more detailed output. This option is helpful for troubleshooting and viewing live progress.

The `rsync` command `-a` or `--archive` option enables "archive mode". This option enables recursive copying and turns on many valuable options to preserve most characteristics of the files. Archive mode is the same as specifying the following options:

Options Enabled with `rsync -a` (Archive Mode)

Option	Description
<code>-r, --recursive</code>	Synchronize the whole directory tree recursively
<code>-l, --links</code>	Synchronize symbolic links
<code>-p, --perms</code>	Preserve permissions
<code>-t, --times</code>	Preserve time stamps
<code>-g, --group</code>	Preserve group ownership
<code>-o, --owner</code>	Preserve the owner of the files
<code>-D, --devices</code>	Preserve device files

Archive mode does not preserve hard links, because it might add significant time to the synchronization. Use the `rsync` command `-H` option to preserve hard links too.

**Note**

To include extended attributes when syncing files, add these options to the `rsync` command:

- `-A` to preserve Access Control Lists (ACLs)
- `-X` to preserve SELinux file contexts

You can use the `rsync` command to synchronize the contents of a local file or directory with a file or directory on a remote machine, with either machine as the source. You can also synchronize the contents of two local files or directories on the same machine.

Like the `sftp` command, the `rsync` command specifies remote locations in the `[user@]host:path` format. The remote location can be either the source or the destination system, provided that one of the two machines is local.

You must be the `root` user on the destination system to preserve file ownership. If the destination is remote, then authenticate as the `root` user. If the destination is local, then you must run the `rsync` command as the `root` user.

In this example, synchronize the local `/var/log` directory to the `/tmp` directory on the `hosta` system:

```
[root@host ~]# rsync -av /var/log hosta:/tmp
root@hosta's password: password
receiving incremental file list
log/
log/README
log/boot.log
...output omitted...
sent 9,783 bytes  received 290,576 bytes  85,816.86 bytes/sec
total size is 11,585,690  speedup is 38.57
```

In the same way, the `/var/log` remote directory on the `hosta` machine synchronizes to the `/tmp` directory on the `host` machine:

```
[root@host ~]# rsync -av hosta:/var/log /tmp
root@hosta's password: password
receiving incremental file list
log/boot.log
log/dnf.librepo.log
log/dnf.log
...output omitted...

sent 9,783 bytes  received 290,576 bytes  85,816.86 bytes/sec
total size is 11,585,690  speedup is 38.57
```

The following example synchronizes the contents of the `/var/log` directory to the `/tmp` directory on the same machine:

```
[user@host ~]$ su -
Password: password
[root@host ~]# rsync -av /var/log /tmp
receiving incremental file list
log/
log/README
log/boot.log
...output omitted...
log/tuned/tuned.log

sent 11,592,423 bytes  received 779 bytes  23,186,404.00 bytes/sec
total size is 11,586,755  speedup is 1.00
[user@host ~]$ ls /tmp
log  ssh-RLjDdarkKiW1
[user@host ~]$
```



Important

Correctly specifying a source directory trailing slash is important. A source directory with a trailing slash synchronizes the contents of the directory *without* including the directory itself. The contents are synced directly into the destination directory. Without the trailing slash, the source directory itself will sync to the destination directory. The source directory's contents are in the new subdirectory in the destination.

Bash Tab-completion automatically adds a trailing slash to directory names.

In this example, the content of the `/var/log/` directory is synchronized to the `/tmp` directory instead of the `log` directory being created in the `/tmp` directory.

```
[root@host ~]# rsync -av /var/log/ /tmp
sending incremental file list
./
README
boot.log
...output omitted...
tuned/tuned.log

sent 11,592,389 bytes  received 778 bytes  23,186,334.00 bytes/sec
total size is 11,586,755  speedup is 1.00
[root@host ~]# ls /tmp
anaconda                dnf.rpm.log-20190318  private
audit                   dnf.rpm.log-20190324  qemu-ga
boot.log                 dnf.rpm.log-20190331  README
...output omitted...
```



References

`rsync(1)` man page

► Guided Exercise

Synchronize Files Between Systems Securely

Synchronize the contents of a local directory with a copy on a remote server, by using the `rsync` command.

Outcomes

- Use the `rsync` command to synchronize the contents of a local directory with a copy on a remote server.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start archive-sync
```

Instructions

- 1. On the `workstation` machine, use the `ssh` command to log in to the `servera` machine as the `student` user, and then switch to the `root` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

- 2. Open a new terminal window, and log in to the `serverb` machine as the `student` user.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
[student@serverb ~]$
```

- 3. Create the `/home/student/serverlogs` directory on the `serverb` machine. Use the `rsync` command to securely create an initial copy of the `/var/log` directory tree on the `servera` machine in the `/home/student/serverlogs` directory on the `serverb` machine.
 - 3.1. On the `serverb` machine, create the `/home/student/serverlogs` directory to store the synchronized log files from the `servera` machine.

```
[student@serverb ~]$ mkdir ~/serverlogs
```

- 3.2. On the `servera` machine, use the `rsync` command to synchronize the `/var/log` directory tree on the `servera` machine to the `/home/student/serverlogs` directory on the `serverb` machine. Only the root user can read all the `/var/log` directory contents on the `servera` machine. Transfer all the files in the initial synchronization.

```
[root@servera ~]# rsync -av /var/log student@serverb:/home/student/serverlogs
...output omitted...
student@serverb's password: student
sending incremental file list
log/
log/README -> ../../usr/share/doc/systemd/README.logs
log/boot.log
...output omitted...

sent 1,390,819 bytes  received 508 bytes  309,183.78 bytes/sec
total size is 1,388,520  speedup is 1.00
```

- ▶ 4. On the `servera` machine, execute the logger `"Log files synchronized"` command to add an entry in the `/var/log/messages` log file to reflect when the last synchronization occurred.

```
[root@servera ~]# logger "Log files synchronized"
```

- ▶ 5. Use the `rsync` command to securely synchronize from the `/var/log` directory tree on the `servera` machine to the `/home/student/serverlogs` directory on the `serverb` machine. This time, only the changed log files are transferred.

```
[root@servera ~]# rsync -av /var/log student@serverb:/home/student/serverlogs
student@serverb's password: student
sending incremental file list
log/messages

sent 3,854 bytes  received 3,807 bytes  2,188.86 bytes/sec
total size is 1,388,648  speedup is 181.26
```

- ▶ 6. On the `serverb` machine, verify the contents of the `/home/student/serverlogs/log/messages` file with the `tail` command.

```
[student@serverb ~]$ tail -n 5 ~/serverlogs/log/messages
Mar 22 04:25:08 servera systemd[1]: systemd-hostnamed.service: Deactivated
successfully.
Mar 22 04:25:09 servera systemd[1066]: Starting Mark boot as successful...
Mar 22 04:25:09 servera systemd[1066]: Finished Mark boot as successful.
Mar 22 04:25:47 servera chronyd[750]: Selected source 172.25.254.254
Mar 22 04:26:25 servera root[1213]: Log files synchronized
```

- ▶ 7. Exit and close the extra terminal.


```
[student@serverb ~]$ exit  
logout  
Connection to serverb closed.  
[student@workstation]$
```

- 8. Return to the `workstation` machine as the `student` user.

```
[student@servera ~]$ exit  
logout  
Connection to servera closed.  
[student@workstation]$
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish archive-sync
```

► Lab

Archive and Transfer Files

Use the `tar`, `rsync`, and `sftp` commands to archive and back up the contents of directories.

Outcomes

- Synchronize a remote directory to a local directory.
- Create an archive of the contents of a synchronized directory.
- Securely copy an archive to a remote host.
- Extract an archive.

Before You Begin

As the `student` user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available. It also installs SSH keys on your systems so that you can transfer files without entering passwords.

```
[student@workstation ~]$ lab start archive-review
```

Instructions

1. On `serverb`, synchronize the `/etc` directory tree from `servera` to the `/conffsync` directory.
2. Create a `configfile-backup-servera.tar.gz` archive with the `/conffsync` directory contents.
3. Securely copy the `/root/configfile-backup-servera.tar.gz` archive file from `serverb` to the `/home/student` directory on `workstation`.
4. On `workstation`, extract the contents to the `/tmp/savedconfig/` directory.

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade archive-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish archive-review
```

► Solution

Archive and Transfer Files

Use the `tar`, `rsync`, and `sftp` commands to archive and back up the contents of directories.

Outcomes

- Synchronize a remote directory to a local directory.
- Create an archive of the contents of a synchronized directory.
- Securely copy an archive to a remote host.
- Extract an archive.

Before You Begin

As the `student` user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available. It also installs SSH keys on your systems so that you can transfer files without entering passwords.

```
[student@workstation ~]$ lab start archive-review
```

Instructions

1. On `serverb`, synchronize the `/etc` directory tree from `servera` to the `/configsync` directory.

- 1.1. Log in to `serverb` as the `student` user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
[student@serverb ~]$ sudo -i
[sudo] password for student: student
[root@serverb ~]#
```

- 1.2. Create the `/configsync` directory to store the synchronized files from `servera`.

```
[root@serverb ~]# mkdir /configsync
```

- 1.3. Synchronize the `/etc` directory tree from `servera` to the `/configsync` directory on `serverb`.

```
[root@serverb ~]# rsync -av root@servera:/etc /configsync
receiving incremental file list
etc/
etc/.pwd.lock
etc/.rht_authorized_keys
...output omitted...
sent 11,519 bytes  received 20,203,285 bytes  40,429,608.00 bytes/sec
total size is 20,150,298  speedup is 1.00
```

2. Create a `configfile-backup-servera.tar.gz` archive with the `/configsync` directory contents.

2.1. Create a `gzip` compressed archive.

```
[root@serverb ~]# tar -czf configfile-backup-servera.tar.gz /configsync
tar: Removing leading `/' from member names
```

2.2. List the contents of the `configfile-backup-servera.tar.gz` archive.

```
[root@serverb ~]# tar -tzf configfile-backup-servera.tar.gz
...output omitted...
configsync/etc/vimrc
configsync/etc/wgetrc
configsync/etc/xattr.conf
```

3. Securely copy the `/root/configfile-backup-servera.tar.gz` archive file from `serverb` to the `/home/student` directory on workstation.

```
[root@serverb ~]# sftp student@workstation
student@workstation's password: student
Connected to workstation.
sftp> put configfile-backup-servera.tar.gz
Uploading configfile-backup-servera.tar.gz to /home/student/configfile-backup-
servera.tar.gz
configfile-backup-servera.tar.gz          100% 4933KB 359.5MB/s   00:00
sftp> bye
```

4. On workstation, extract the contents to the `/tmp/savedconfig/` directory.

4.1. Return to the workstation system as the `student` user.

```
[root@serverb ~]# exit
logout
[student@serverb ~]$ exit
logout
Connection to serverb closed.
[student@workstation]$
```

4.2. Create the `/tmp/savedconfig` directory, to store the extracted contents. Change to the new directory.

```
[student@workstation ~]$ mkdir /tmp/savedconfig
[student@workstation ~]$ cd /tmp/savedconfig
[student@workstation savedconfig]$
```

4.3. List the contents of the `configfile-backup-servera.tar.gz` archive.

```
[student@workstation savedconfig]$ tar -tzf ~/configfile-backup-servera.tar.gz
...output omitted...
configsync/etc/vimrc
configsync/etc/wgetrc
configsync/etc/xattr.conf
```

4.4. Extract the archive contents to the `/tmp/savedconfig/` directory.

```
[student@workstation savedconfig]$ tar -xzf ~/configfile-backup-servera.tar.gz
```

4.5. List the directory to verify that it contains `etc` files.

```
[student@workstation savedconfig]$ ls -lR
.:
total 0
drwxr-xr-x. 3 student student 17 Mar 28 16:32 configsync

./configsync:
total 12
drwxr-xr-x. 105 student student 8192 Mar 28 16:03 etc
...output omitted...
```

4.6. Return to the student user's home directory.

```
[student@workstation savedconfig]$ cd
[student@workstation ~]$
```

Evaluation

As the student user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade archive-review
```

Finish

On the `workstation` machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish archive-review
```

Summary

- The `tar` command creates an archive file from a set of files and directories. This command also extracts and lists files from an archive file.
- The `tar` command provides a set of compression methods to reduce archive size.
- Besides providing a secure remote shell, the SSH service also provides the `sftp` command to transfer files securely to and from a remote system that runs the SSH server.
- The `rsync` command securely and efficiently synchronizes files between two directories, of which either one can be on a remote system.

Chapter 5

Tune System Performance

Goal

Improve system performance by setting tuning parameters and adjusting the scheduling priority of processes.

Objectives

- Optimize system performance by selecting a tuning profile that the `tuned` daemon manages.
- Prioritize or deprioritize specific processes, with the `nice` and `renice` commands.

Sections

- Adjust Tuning Profiles (and Guided Exercise)
- Influence Process Scheduling (and Guided Exercise)

Lab

- Tune System Performance

Adjust Tuning Profiles

Objectives

- Optimize system performance by selecting a tuning profile that the `tuned` daemon manages.

Tune Systems

System administrators optimize the performance of a system by adjusting device settings based on various use case workloads. The `tuned` daemon applies tuning adjustments both statically and dynamically by using tuning profiles that reflect particular workload requirements.

Configure Static Tuning

The `tuned` daemon applies system settings when a service starts or on selecting a new tuning profile. Static tuning configures predefined `kernel` parameters in profiles that the `tuned` daemon applies at boot. With static tuning, the `tuned` daemon sets kernel parameters for overall performance expectations, without adjusting these parameters as activity levels change.

Configure Dynamic Tuning

With dynamic tuning, the `tuned` daemon monitors system activity and adjusts settings according to runtime behavior changes. Dynamic tuning continuously adjusts tuning to fit the current workload, starting with the initial declared settings in your selected tuning profile.

For example, storage devices experience high use during startup and login, but have minimal activity when user workloads consist of using web browsers and email clients. Similarly, CPU and network devices experience activity increases during peak usage throughout a workday. The `tuned` daemon monitors the activity of these components, and adjusts parameter settings to maximize performance during high-activity times and to reduce settings during low activity. Predefined tuning profiles provide performance parameters that the `tuned` daemon uses.

To monitor and adjust parameter settings, the `tuned` daemon uses modules called monitor and tuning plug-ins, respectively.

Monitor plug-ins analyze the system and obtain information from it, so the tuning plug-ins use this information for dynamic tuning. At this moment, the `tuned` daemon ships with three monitor plug-ins:

- `disk`: Monitors the disk load based on the number of I/O operations for every disk device.
- `net`: Monitors the network load based on the number of transferred packets per network card.
- `load`: Monitors the CPU load for every CPU.

Tuning plug-ins tune the individual subsystems. They use the data from the monitor plug-ins and the performance parameters from the predefined tuning profiles. Among others, the `tuned` daemon ships with the following tuning plug-ins:

- `disk`: Sets different disk parameters, for example, the disk scheduler, the spin-down timeout, or the advanced power management.
- `net`: Configures the interface speed and the Wake-on-LAN (WoL) functionality.
- `cpu`: Sets different CPU parameters, for example, the CPU governor or the latency.

By default, dynamic tuning is disabled. You can enable it by setting the `dynamic_tuning` variable to 1 in the `/etc/tuned/tuned-main.conf` configuration file. If you enable dynamic tuning, then the `tuned` daemon periodically monitors the system and adjusts the tuning settings to runtime behavior changes. You can set the time in seconds between updates by using the `update_interval` variable in the `/etc/tuned/tuned-main.conf` configuration file.

```
[root@host ~]$ cat /etc/tuned/tuned-main.conf
...output omitted...
# Dynamically tune devices, if disabled only static tuning will be used.
dynamic_tuning = 1
...output omitted...
# Update interval for dynamic tunings (in seconds).
# It must be multiply of the sleep_interval.
update_interval = 10
...output omitted...
```

The tuned Utility

You can install and enable the `tuned` package by using the following commands:

```
[root@host ~]$ dnf install tuned
...output omitted...
[root@host ~]$ systemctl enable --now tuned
Created symlink /etc/systemd/system/multi-user.target.wants/tuned.service → /usr/lib/systemd/system/tuned.service.
```

The `tuned` application provides profiles in the following categories:

- Power-saving profiles
- Performance-boosting profiles

The performance-boosting profiles include profiles that focus on the following aspects:

- Low latency for storage and network
- High throughput for storage and network
- Virtual machine performance
- Virtualization host performance

The next table shows a list of the tuning profiles that are distributed with Red Hat Enterprise Linux 9:

Tuning Profiles Distributed with Red Hat Enterprise Linux 9

Tuned Profile	Purpose
balanced	Ideal for systems that require a compromise between power saving and performance.
powersave	Tunes the system for maximum power saving.
throughput-performance	Tunes the system for maximum throughput.
accelerator-performance	Tunes the same as <code>throughput-performance</code> , and also reduces the latency to less than 100 μ s.

Tuned Profile	Purpose
latency-performance	Ideal for server systems that require low latency at the expense of power consumption.
network-throughput	Derived from the throughput-performance profile. Additional network tuning parameters are applied for maximum network throughput.
network-latency	Derived from the latency-performance profile. Enables additional network tuning parameters to provide low network latency.
desktop	Derived from the balanced profile. Provides faster response of interactive applications.
hpc-compute	Derived from the latency-performance profile. Ideal for high-performance computing.
virtual-guest	Tunes the system for maximum performance if it runs on a virtual machine.
virtual-host	Tunes the system for maximum performance if it acts as a host for virtual machines.
intel-sst	Optimized for systems with Intel Speed Select Technology configurations. Use it as an overlay on other profiles.
optimize-serial-console	Increases responsiveness of the serial console. Use it as an overlay on other profiles.

The `tuned` application stores the tuning profiles under the `/usr/lib/tuned` and `/etc/tuned` directories. Every profile has a separate directory, and inside the directory the `tuned.conf` main configuration file and, optionally, other files.

```
[root@host ~]# cd /usr/lib/tuned
[root@host tuned]# ls
accelerator-performance  hpc-compute          network-throughput    throughput-
performance
balanced                  intel-sst             optimize-serial-console virtual-
guest
desktop                   latency-performance  powersave             virtual-
host
functions                  network-latency      recommend.d
[root@host tuned]$ ls virtual-guest
tuned.conf
```

A typical `tuned.conf` configuration file looks as follows:

```
[root@host tuned]# cat virtual-guest/tuned.conf
#
# tuned configuration
#
```

```
[main]
summary=Optimize for running inside a virtual guest
include=throughput-performance

[sysctl]
# If a workload mostly uses anonymous memory and it hits this limit, the entire
# working set is buffered for I/O, and any more write buffering would require
# swapping, so it's time to throttle writes until I/O can catch up. Workloads
# that mostly use file mappings may be able to use even higher values.
#
# The generator of dirty data starts writeback at this percentage (system default
# is 20%)
vm.dirty_ratio = 30

# Filesystem I/O is usually much more efficient than swapping, so try to keep
# swapping low. It's usually safe to go even lower than this on systems with
# server-grade storage.
vm.swappiness = 30
```

The `[main]` section in the file might include a summary of the tuning profile. This section also accepts the `include` parameter, for the profile to inherit all the settings from the referenced profile.

This configuration file is useful when creating new tuning profiles, because you can use one of the provided profiles as a basis, and then add or modify the parameters to configure. To create or modify tuning profiles, copy the tuning profile files from the `/usr/lib/tuned` directory to the `/etc/tuned` directory and then modify them. If profile directories exist with the same name under the `/usr/lib/tuned` and `/etc/tuned` directories, the latter always take precedence. Thus, never directly modify files in the `/usr/lib/tuned` system directory.

The rest of the sections in the `tuned.conf` file use the tuning plug-ins to modify parameters in the system. In the previous example, the `[sysctl]` section modifies the `vm.dirty_ratio` and `vm.swappiness` kernel parameters through the `sysctl` plug-in.

Manage Profiles from the Command Line

Use the `tuned-adm` command to change the settings of the `tuned` daemon. The `tuned-adm` command queries current settings, lists available profiles, recommends a tuning profile for the system, changes profiles directly, or turns off tuning.

You can identify the currently active tuning profile with the `tuned-adm active` command.

```
[root@host ~]# tuned-adm active
Current active profile: virtual-guest
```

The `tuned-adm list` command lists all available tuning profiles, including both built-in profiles and the custom-created tuning profiles.

```
[root@host ~]# tuned-adm list
Available profiles:
- accelerator-performance    - Throughput performance based tuning with ...
- balanced                   - General non-specialized tuned profile
- desktop                    - Optimize for the desktop use-case
- hpc-compute                 - Optimize for HPC compute workloads
```

```
- intel-sst           - Configure for Intel Speed Select Base Frequency
- latency-performance - Optimize for deterministic performance at ...
- network-latency    - Optimize for deterministic performance at ...
...output omitted...
Current active profile: virtual-guest
```

Use the `tuned-adm profile_info` command for information about a given profile.

```
[root@host ~]$ tuned-adm profile_info network-latency
Profile name:
network-latency

Profile summary:
Optimize for deterministic performance at the cost of increased power consumption,
  focused on low latency network performance
...output omitted..
```

If no profile is specified, then the `tuned-adm profile_info` command shows the information for the active tuning profile:

```
[root@host ~]$ tuned-adm active
Current active profile: virtual-guest
[root@host ~]$ tuned-adm profile_info
Profile name:
virtual-guest

Profile summary:
Optimize for running inside a virtual guest
...output omitted..
```

Use the `tuned-adm profile profilename` command to switch to a different active profile that better matches the system's current tuning requirements.

```
[root@host ~]$ tuned-adm profile throughput-performance
[root@host ~]$ tuned-adm active
Current active profile: throughput-performance
```

The `tuned-adm recommend` command can recommend a tuning profile for the system. The system uses this mechanism to determine the default profile after its installation.

```
[root@host ~]$ tuned-adm recommend
virtual-guest
```

Note

The `tuned-adm recommend` command bases its recommendation on various system characteristics, including whether the system is a virtual machine and other predefined selected categories during system installation.

To revert the setting changes that the current profile applied, either switch to another profile or deactivate the tuned daemon. Turn off the tuned application tuning activity by using the `tuned-adm off` command.

```
[root@host ~]$ tuned-adm off
[root@host ~]$ tuned-adm active
No current active profile.
```

Manage Profiles with the Web Console

To manage system performance profiles with the web console, you must log in and escalate privileges. Privilege escalation mode permits the user to execute commands, with administrative privileges, that modify system performance profiles. Because changing tuning profiles modifies some system parameters, you must do it with administrative privileges.

You can switch to the administrative access mode in the web console by clicking the **Limited access** or the **Turn on administrative access** buttons. Then, enter your password when prompted. After you escalate privileges, the **Limited access** button changes to **Administrative access**. As a security reminder, always toggle back to limited access mode after completing the system task that requires administrative privileges.

As a privileged user, click the **Overview** menu option in the left navigation bar. The **Performance profile** field displays the current active profile.




Figure 5.1: Active performance profile

To select a different profile, click the active profile link. In the **Change performance profile** user interface, scroll through the profile list to select one that best suits the system purpose, and click the **Change profile** button.




Figure 5.2: Select a preferred performance profile

To verify changes, return to the main **Overview** page, and confirm that it displays the active profile in the **Performance profile** field.



References

`tuned(8)`, `tuned.conf(5)`, `tuned-main.conf(5)`, and `tuned-adm(1)` man pages

For more information, refer to the *Monitoring and Managing System Status and Performance* guide at https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/monitoring_and_managing_system_status_and_performance/index

► Guided Exercise

Adjust Tuning Profiles

Tune server performance by activating the tuned service and applying a tuning profile.

Outcomes

- Configure a system to use a tuning profile.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start tuning-profiles
```

Instructions

- 1. Log in to `servera` as the `student` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
```

- 2. Verify that the `tuned` package is installed, enabled, and started.

- 2.1. Verify that the `tuned` package is installed.

```
[student@servera ~]$ dnf list tuned
...output omitted...
Installed Packages
tuned.noarch                2.18.0-1.el9                @System
```

- 2.2. Verify that the service is enabled.

```
[student@servera ~]$ systemctl is-enabled tuned
enabled
```

- 2.3. Verify that the service is currently running.

```
[student@servera ~]$ systemctl is-active tuned
active
```

- 3. List the available tuning profiles and identify the active profile.


```
[student@servera ~]$ sudo tuned-adm list
[sudo] password for student: student
Available profiles:
- accelerator-performance      - Throughput performance based tuning with disabled
  higher latency STOP states
- balanced                     - General non-specialized tuned profile
- desktop                      - Optimize for the desktop use-case
- hpc-compute                  - Optimize for HPC compute workloads
- intel-sst                    - Configure for Intel Speed Select Base Frequency
- latency-performance          - Optimize for deterministic performance at the cost
  of increased power consumption
- network-latency              - Optimize for deterministic performance at the cost
  of increased power consumption, focused on low latency network performance
- network-throughput           - Optimize for streaming network throughput,
  generally only necessary on older CPUs or 40G+ networks
- optimize-serial-console       - Optimize for serial console use.
- powersave                   - Optimize for low power consumption
- throughput-performance       - Broadly applicable tuning that provides excellent
  performance across a variety of common server workloads
- virtual-guest                - Optimize for running inside a virtual guest
- virtual-host                 - Optimize for running KVM guests
Current active profile: virtual-guest
```

- ▶ 4. Review the `tuned.conf` configuration file for the current active profile, `virtual-guest`. You can find the `tuned.conf` configuration file in the `/usr/lib/tuned/virtual-guest` directory. The `virtual-guest` tuning profile is based on the `throughput-performance` profile, but it sets different values for the `vm.dirty_ratio` and `vm.swappiness` parameters. Verify that the `virtual-guest` tuning profile applies these values on your system.
 - 4.1. Review the `virtual-guest` configuration file in the `/usr/lib/tuned/virtual-guest` directory. Verify the values for the `vm.dirty_ratio` and `vm.swappiness` parameters.

```
[student@servera ~]$ cat /usr/lib/tuned/virtual-guest/tuned.conf
#
# tuned configuration
#

[main]
summary=Optimize for running inside a virtual guest
include=throughput-performance

[sysctl]
# If a workload mostly uses anonymous memory and it hits this limit, the entire
# working set is buffered for I/O, and any more write buffering would require
# swapping, so it's time to throttle writes until I/O can catch up. Workloads
# that mostly use file mappings may be able to use even higher values.
#
# The generator of dirty data starts writeback at this percentage (system default
# is 20%)
vm.dirty_ratio = 30
```

```
# Filesystem I/O is usually much more efficient than swapping, so try to keep
# swapping low. It's usually safe to go even lower than this on systems with
# server-grade storage.
vm.swappiness = 30
```

4.2. Verify that the tuning profile applies these values on your system.

```
[student@servera ~]$ sysctl vm.dirty_ratio
vm.dirty_ratio = 30
[student@servera ~]$ sysctl vm.swappiness
vm.swappiness = 30
```

- ▶ 5. Review the `tuned.conf` configuration file for the `virtual-guest` parent's tuning profile, `throughput-performance`. You can find it in the `/usr/lib/tuned/throughput-performance` directory. Notice that the `throughput-performance` tuning profile sets a different value for the `vm.dirty_ratio` and `vm.swappiness` parameters, although the `virtual-guest` profile overwrites them. Verify that the `virtual-guest` tuning profile applies the value for the `vm.dirty_background_ratio` parameter, which it inherits from the `throughput-performance` profile.

5.1. Review the `throughput-performance` configuration file in the `/usr/lib/tuned/throughput-performance` directory. Verify the values for the `vm.dirty_ratio`, `vm.swappiness`, and `vm.dirty_background_ratio` parameters.

```
[student@servera ~]$ cat /usr/lib/tuned/throughput-performance/tuned.conf
#
# tuned configuration
#

[main]
summary=Broadly applicable tuning that provides excellent performance across a
variety of common server workloads

...output omitted...

[sysctl]
# If a workload mostly uses anonymous memory and it hits this limit, the entire
# working set is buffered for I/O, and any more write buffering would require
# swapping, so it's time to throttle writes until I/O can catch up. Workloads
# that mostly use file mappings may be able to use even higher values.
#
# The generator of dirty data starts writeback at this percentage (system default
# is 20%)
vm.dirty_ratio = 40

# Start background writeback (via writeback threads) at this percentage (system
# default is 10%)
vm.dirty_background_ratio = 10

# PID allocation wrap value. When the kernel's next PID value
# reaches this value, it wraps back to a minimum PID value.
# PIDs of value pid_max or larger are not allocated.
#
```

```
# A suggested value for pid_max is 1024 * <# of cpu cores/threads in system>
# e.g., a box with 32 cpus, the default of 32768 is reasonable, for 64 cpus,
# 65536, for 4096 cpus, 4194304 (which is the upper limit possible).
#kernel.pid_max = 65536

# The swappiness parameter controls the tendency of the kernel to move
# processes out of physical memory and onto the swap disk.
# 0 tells the kernel to avoid swapping processes out of physical memory
# for as long as possible
# 100 tells the kernel to aggressively swap processes out of physical memory
# and move them to swap cache
vm.swappiness=10

...output omitted...
```

- 5.2. Verify that the `virtual-guest` tuning profile applies the inherited `vm.dirty_background_ratio` parameter.

```
[student@servera ~]$ sysctl vm.dirty_background_ratio
vm.dirty_background_ratio = 10
```

- ▶ 6. Change the current active tuning profile to `throughput-performance`, and then confirm the results. Verify that the `vm.dirty_ratio` and `vm.swappiness` parameters change to the values in the `throughput-performance` configuration file.

- 6.1. Change the current active tuning profile.

```
[student@servera ~]$ sudo tuned-adm profile throughput-performance
```

- 6.2. Confirm that `throughput-performance` is the active tuning profile.

```
[student@servera ~]$ sudo tuned-adm active
Current active profile: throughput-performance
```

- 6.3. Verify the values for the `vm.dirty_ratio` and `vm.swappiness` parameters.

```
[student@servera ~]$ sysctl vm.dirty_ratio
vm.dirty_ratio = 40
[student@servera ~]$ sysctl vm.swappiness
vm.swappiness = 10
```

- ▶ 7. Return to the workstation machine as the student user.

```
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish tuning-profiles
```

Influence Process Scheduling

Objectives

- Prioritize or deprioritize specific processes, with the `nice` and `renice` commands.

Linux Process Scheduling

Modern computer systems use multi-core, multi-thread CPUs that can execute many instruction threads simultaneously. The largest high-performing supercomputers can have hundreds or thousands of CPUs with hundreds of processing cores and thread structures per CPU, and can process millions of instruction threads in parallel. Although a single user who runs many applications can saturate the typical desktop system or personal workstation with CPU activity, a correctly sized and configured workstation is designed to match the user's intended workload. However, the typical enterprise or internet server handles many hundreds or thousands of users and application requests each second, which can result in CPU saturation. All systems under CPU load experience scenarios that require handling more process threads than the available system CPU processing units to schedule the threads immediately.

Linux and other operating systems use a technique called *time-slicing* or *multitasking* for process management. The operating system *process scheduler* rapidly switches between process threads on each available CPU core. This behavior gives the impression that many processes are running at the same time.

Process Priorities

Process *priority* sets the importance of each process. Linux implements *scheduling policies* that define the rules by which processes are organized and prioritized to obtain CPU processing time. The various Linux scheduling policies might be designed to handle interactive application requests, or non-interactive batch application processing, or real-time application requirements. Real-time scheduling policies still use process priorities and queues. However, current, non-real-time (*normal*) scheduling policies use the Completely Fair Scheduler (CFS), which instead organizes processes that are awaiting CPU time into a binary search tree. This process priority uses the `SCHED_NORMAL` or `SCHED_OTHER` policies as the default scheduling policy.

Processes that run under the `SCHED_NORMAL` policy are assigned a *static* real-time priority of 0, to ensure that all system real-time processes have a higher priority than normal processes. However, this static priority value is not included when organizing normal process threads for CPU scheduling. Instead, the CFS scheduling algorithm arranges normal process threads into a time-weighted binary tree, where the first item has the lowest previously spent CPU time, and the last item has the most cumulative CPU time.

Nice Value

The order of the binary tree is additionally influenced by a user-modifiable, per-process *nice* value, which ranges from -20 (increased priority) to 19 (decreased priority), with a default of 0. Processes inherit their starting nice value from their parent process. All users can adjust the nice value to decrease priority, but only `root` can increase its priority.

A higher nice value indicates a decrease in the process priority from the default, or *making the process nicer* to other processes. A lower nice value indicates an increase in the process priority from the default, or *making the process less nice* to other processes.

Increasing the nice value lowers the thread's position, and decreasing the value raises the thread's position.



Important

Generally, priorities determine only indirectly the amount of CPU time that a process thread receives. On a non-saturated system with available CPU capacity, every process is scheduled for immediate CPU time, for as much time as each process wants. Relative process importance, as managed in the binary tree, determines only which threads are selected and placed on CPUs first.

On a CPU-saturated system, where more waiting threads than CPU processing units exist, higher-priority (lower nice) process threads are placed first, until all CPU units are busy, while the lower-priority (higher nice) threads initially must wait in the binary tree. However, the Completely Fair Scheduler is designed to balance process importance, nice values, and previous cumulative CPU time, and dynamically adjusts the binary tree such that all processes obtain fair CPU time.

Permission to Modify Nice Values

Privileged users can decrease the nice value of a process, to make a process less nice. A process is then repetitively placed higher in the binary tree, and therefore is scheduled more often. On a saturated system, the overall CPU time available to other processes is reduced.

Unprivileged users can only increase the nice value on their own processes, which makes their own processes nicer, and therefore lowers their placement in the binary tree. Unprivileged users cannot decrease their processes' nice values to raise their importance, nor can they adjust the nice values for another user's process.

Viewing Nice Values

Nice values map to a priority value, and both values are available for viewing in process listing commands. A nice value of -20 maps to a 0 priority in the `top` command. A nice value of 19 maps to a 39 priority in the `top` command.




Figure 5.3: Priorities and nice values as reported by the `top` command

In the preceding figure, the nice values are aligned with the priority values that are used by the `top` command. The `top` command displays the process priority in the PR column, and the nice value in the NI column. The `top` priority numbering scheme, which displays real-time process priorities as negative numbers, is a legacy of internal priority algorithms.

The following output is the summary and a partial process listing in the `top` command:

```
Tasks: 192 total,   1 running, 191 sleeping,   0 stopped,   0 zombie
%Cpu(s):  0.0 us,  1.6 sy,   0.0 ni, 96.9 id,   0.0 wa,   0.0 hi,   1.6 si,   0.0 st
MiB Mem :  5668.6 total,  4655.6 free,   470.1 used,   542.9 buff/cache
MiB Swap:    0.0 total,    0.0 free,    0.0 used.  4942.6 avail Mem

  PID USER      PR  NI   VIRT   RES   SHR S  %CPU  %MEM    TIME+  COMMAND
    1 root        20   0  172180  16232  10328 S   0.0   0.3   0:01.49 systemd
    2 root        20   0        0     0      0 S   0.0   0.0   0:00.01 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
```

The `ps` command displays process nice values, when using the default formatting options.

The following `ps` command lists all processes with their process ID, process name, nice value, and scheduling class. The processes are sorted in descending order by nice value. In the CLS scheduling class column, TS stands for *time sharing*, which is another name for the normal scheduling policies, including `SCHED_NORMAL`. Other CLS values, such as FF for *first in first out* and RR for *round robin*, indicate real-time processes. Real-time processes are not assigned nice values, as indicated by the dash (-) in the NI column. Advanced scheduling policies are taught in the *Red Hat Performance Tuning: Linux in Physical, Virtual, and Cloud (RH442)* course.

```
[user@host ~]$ ps axo pid,comm,nice,cls --sort=-nice
PID COMMAND          NI CLS
  33 khugepaged        19 TS
  32 ksm               5 TS
 814 rtkit-daemon       1 TS
   1 systemd           0 TS
   2 kthreadd           0 TS
   5 kworker/0:0-cgr     0 TS
   7 kworker/0:1-rcu     0 TS
   8 kworker/u4:0-ev     0 TS
  15 migration/0       -  FF
...output omitted...
```

Start Processes with User-set Nice Values

When a process is created, it inherits its parent's nice value. When a process starts from the command line, it inherits its nice value from the shell process. Typically, new processes run with the default nice value of 0.

The following example starts a process from the shell, and displays the process's nice value. Note the use of the PID option in the `ps` command to specify the requested output.

Note

This command was chosen for demonstration for its low resource consumption.

```
[user@host ~]$ sleep 60 &
[1] 2667
[user@host ~]$ ps -o pid,comm,nice 2667
  PID COMMAND      NI
 2667 sleep         0
```

All users can use the `nice` command to start commands with a default or higher nice value. Setting a higher value by default ensures that the new process is a lower priority than your current working shell, and would less likely affect your current interactive session.

The following example starts the same command as a background job with the default nice value, and displays the process's nice value:

```
[user@host ~]$ nice sleep 60 &
[1] 2736
[user@host ~]$ ps -o pid,comm,nice 2736
  PID COMMAND      NI
 2736 sleep        10
```

Use the `nice` command `-n` option to apply a user-defined nice value to the starting process. The default is to add 10 to the process's current nice value. The following example starts a background job with a user-defined nice value of 15 and displays the result:

```
[user@host ~]$ nice -n 15 sleep 60 &
[1] 2673
[user@host ~]$ ps -o pid,comm,nice 2740
  PID COMMAND      NI
 2740 sleep        15
```

Change the Nice Value of an Existing Process

You can change the nice value of an existing process with the `renice` command. This example uses the process ID from the previous example to change from the current nice value of 15 to a new nice value of 19.

```
[user@host ~]$ renice -n 19 2740
2740 (process ID) old priority 15, new priority 19
```

You can also use the `top` command to change the nice value on an existing process. From the `top` interactive interface, press the `r` key to access the `renice` command. Enter the process ID, and then enter the new nice value.



References

`nice(1)`, `renice(1)`, `top(1)`, and `sched_setscheduler(2)` man pages

► Guided Exercise

Influence Process Scheduling

Adjust the scheduling priority of processes with the `nice` and `renice` commands, and observe the effects on process execution.

Outcomes

- Adjust scheduling priorities for processes.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command ensures that all required resources are available.

```
[student@workstation ~]$ lab start tuning-procscheduling
```



Important

This exercise uses commands that perform an endless checksum on a device file and intentionally use significant CPU resources.

Instructions

- 1. Use the `ssh` command to log in to the `servera` machine as the `student` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
```

- 2. Determine the number of CPU cores on the `servera` machine, and then start two instances of the `sha1sum /dev/zero &` command for each core.
 - 2.1. Use the `grep` command to parse the number of existing virtual processors (CPU cores) from the `/proc/cpuinfo` file.

```
[student@servera ~]$ grep -c '^processor' /proc/cpuinfo
2
```

- 2.2. Use a looping command to start multiple instances of the `sha1sum /dev/zero &` command. Start two instances for each virtual processor that was indicated in the previous step. In this example, a `for` loop creates four instances. The PID values in your output might vary from the example.

```
[student@servera ~]$ for i in {1..4}; do sha1sum /dev/zero & done
[1] 1132
[2] 1133
[3] 1134
[4] 1135
```

- ▶ 3. Verify that the background jobs are running for each of the `sha1sum` processes.

```
[student@servera ~]$ jobs
[1]   Running                  sha1sum /dev/zero &
[2]   Running                  sha1sum /dev/zero &
[3]-  Running                  sha1sum /dev/zero &
[4]+  Running                  sha1sum /dev/zero &
```

- ▶ 4. Use the `ps` and `pgrep` commands to display the percentage of CPU usage for each `sha1sum` process.

```
[student@servera ~]$ ps u $(pgrep sha1sum)
USER      PID %CPU %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
student  1132 49.6   0.1 225336  2288 pts/0    R   11:40   2:40 sha1sum /dev/zero
student  1133 49.6   0.1 225336  2296 pts/0    R   11:40   2:40 sha1sum /dev/zero
student  1134 49.6   0.1 225336  2264 pts/0    R   11:40   2:40 sha1sum /dev/zero
student  1135 49.6   0.1 225336  2280 pts/0    R   11:40   2:40 sha1sum /dev/zero
```

- ▶ 5. Terminate all `sha1sum` processes, and then verify that no jobs are running.

- 5.1. Use the `pkill` command to terminate all running processes with the `sha1sum` name pattern.

```
[student@servera ~]$ pkill sha1sum
[2]   Terminated              sha1sum /dev/zero
[4]+  Terminated              sha1sum /dev/zero
[1]-  Terminated              sha1sum /dev/zero
[3]+  Terminated              sha1sum /dev/zero
```

- 5.2. Verify that no jobs are running.

```
[student@servera ~]$ jobs
[student@servera ~]$
```

- ▶ 6. Start multiple instances of the `sha1sum /dev/zero &` command, and then start one additional instance of the `sha1sum /dev/zero &` command with a `nice` level of 10. Start at least as many instances as the number of system virtual processors. In this example, three regular instances are started, plus another with a higher `nice` level.

- 6.1. Use looping to start three instances of the `sha1sum /dev/zero &` command.

```
[student@servera ~]$ for i in {1..3}; do sha1sum /dev/zero & done
[1] 1207
[2] 1208
[3] 1209
```

6.2. Use the `nice` command to start the fourth instance with a nice level of 10.

```
[student@servera ~]$ nice -n 10 sha1sum /dev/zero &
[4] 1210
```

- ▶ 7. Use the `ps` and `pgrep` commands to display the PID, percentage of CPU usage, nice value, and executable name for each process. The instance with the nice value of 10 displays a lower percentage of CPU usage than the other instances.

```
[student@servera ~]$ ps -o pid,pcpu,nice,comm $(pgrep sha1sum)
  PID %CPU  NI  COMMAND
 1207  64.2   0  sha1sum
 1208  65.0   0  sha1sum
 1209  63.9   0  sha1sum
 1210  8.2  10  sha1sum
```

- ▶ 8. Use the `sudo renice` command to lower the nice level of a process from the previous step. Use the PID value of the process instance with the nice level of 10 to lower its nice level to 5.

```
[student@servera ~]$ sudo renice -n 5 1210
[sudo] password for student: student
1210 (process ID) old priority 10, new priority 5
```

- ▶ 9. Repeat the `ps` and `pgrep` commands to display the CPU percentage and nice level.

```
[student@servera ~]$ ps -o pid,pcpu,nice,comm $(pgrep sha1sum)
  PID %CPU  NI  COMMAND
 1207  62.9   0  sha1sum
 1208  63.2   0  sha1sum
 1209  63.2   0  sha1sum
 1210 10.9   5  sha1sum
```

- ▶ 10. Use the `pkill` command to terminate all running processes with the `sha1sum` name pattern.

```
[student@servera ~]$ pkill sha1sum
...output omitted...
```

- ▶ 11. Return to the `workstation` machine as the `student` user.

```
[student@servera ~]$ exit  
logout  
Connection to servera closed.  
[student@workstation ~]$
```



Important

Verify that you have terminated all exercise processes before leaving this exercise.

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish tuning-procscheduling
```

► Lab

Tune System Performance

Apply a specific tuning profile and adjust the scheduling priority of an existing process with high CPU usage.

Outcomes

- Activate a specific tuning profile for a computer system.
- Adjust the CPU scheduling priority of a process.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start tuning-review
```



Important

This lab uses commands that perform an endless checksum on a device file and intentionally use significant CPU resources.

Instructions

1. Change the current tuning profile for the `serverb` machine to the `balanced` profile, a general non-specialized tuned profile. List the information for the `balanced` tuning profile when it is the current tuning profile.
2. Two processes on `serverb` are consuming a high percentage of CPU usage. Adjust each process's `nice` level to 10 to allow more CPU time for other processes.

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade tuning-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish tuning-review
```

► Solution

Tune System Performance

Apply a specific tuning profile and adjust the scheduling priority of an existing process with high CPU usage.

Outcomes

- Activate a specific tuning profile for a computer system.
- Adjust the CPU scheduling priority of a process.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start tuning-review
```



Important

This lab uses commands that perform an endless checksum on a device file and intentionally use significant CPU resources.

Instructions

1. Change the current tuning profile for the `serverb` machine to the `balanced` profile, a general non-specialized tuned profile. List the information for the `balanced` tuning profile when it is the current tuning profile.

- 1.1. Log in to the `serverb` machine as the `student` user.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
[student@serverb ~]$
```

- 1.2. Verify that the `tuned` package is installed.

```
[student@serverb ~]$ dnf list tuned
...output omitted...
Installed Packages
tuned.noarch                2.18.0-1.el9                @System
```

- 1.3. Verify the `tuned` service state.

```
[student@serverb ~]$ systemctl is-active tuned
active
```

- 1.4. List all available tuning profiles and their descriptions. Note that the current active profile is **virtual-guest**.

```
[student@serverb ~]$ sudo tuned-adm list
[sudo] password for student: student
Available profiles:
- accelerator-performance - Throughput performance based tuning with disabled
                           higher latency STOP states
- balanced                 - General non-specialized tuned profile
- desktop                 - Optimize for the desktop use-case
- hpc-compute             - Optimize for HPC compute workloads
- intel-sst               - Configure for Intel Speed Select Base Frequency
- latency-performance     - Optimize for deterministic performance at the cost
                           of increased power consumption
- network-latency         - Optimize for deterministic performance at the cost
                           of increased power consumption, focused on low
                           latency network performance
- network-throughput      - Optimize for streaming network throughput, generally
                           only necessary on older CPUs or 40G+ networks
- optimize-serial-console - Optimize for serial console use.
- powersave              - Optimize for low power consumption
- throughput-performance - Broadly applicable tuning that provides excellent
                           performance across a variety of common server
                           workloads
- virtual-guest           - Optimize for running inside a virtual guest
- virtual-host            - Optimize for running KVM guests
Current active profile: virtual-guest
```

- 1.5. Change the current active tuning profile to the **balanced** profile.

```
[student@serverb ~]$ sudo tuned-adm profile balanced
```

- 1.6. List summary information of the current active tuned profile. Verify that the active profile is the **balanced** profile.

```
[student@serverb ~]$ sudo tuned-adm profile_info
Profile name:
balanced

Profile summary:
General non-specialized tuned profile
...output omitted...
```

2. Two processes on **serverb** are consuming a high percentage of CPU usage. Adjust each process's **nice** level to 10 to allow more CPU time for other processes.
 - 2.1. Determine the top two CPU consumers on the **serverb** machine. The **ps** command lists the top CPU consumers at the bottom of the output. CPU percentage values might vary on your machine.


```
[student@serverb ~]$ ps aux --sort=pcpu
USER      PID %CPU %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
...output omitted...
root      1079 98.5  0.1 225340 2300 ?        RN    06:25   4:29 sha1sum /dev/zero
root      1095 99.0  0.1 225340 2232 ?        R<    06:25   4:30 md5sum /dev/zero
```

2.2. Identify the current `nice` level for each of the top two CPU consumers.

```
[student@serverb ~]$ ps -o pid,pcpu,nice,comm \
$(pgrep sha1sum;pgrep md5sum)
  PID %CPU  NI COMMAND
 1079  98.8   2 sha1sum
 1095  99.1  -2 md5sum
```

2.3. Adjust the `nice` level for each process to 10. Use the correct PID values for your processes from the previous command output.

```
[student@serverb ~]$ sudo renice -n 10 1079 1095
[sudo] password for student: student
1079 (process ID) old priority 2, new priority 10
1095 (process ID) old priority -2, new priority 10
```

2.4. Verify that the current `nice` level for each process is 10.

```
[student@serverb ~]$ ps -o pid,pcpu,nice,comm \
$(pgrep sha1sum;pgrep md5sum)
  PID %CPU  NI COMMAND
 1079  98.9  10 sha1sum
 1095  99.2  10 md5sum
```

2.5. Return to the `workstation` machine as the `student` user.

```
[student@serverb ~]$ exit
logout
Connection to serverb closed.
[student@workstation ~]$
```



Important

Verify that you have terminated all lab processes before leaving this lab.

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade tuning-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish tuning-review
```

Summary

- The `tuned` service automatically modifies device settings to meet specific system needs based on a predefined selected tuning profile.
- To revert all changes of the selected profile to the system settings, either switch to another profile or deactivate the `tuned` service.
- The system assigns a relative priority to a process to determine its CPU access. This priority is called the `nice` value of a process.
- The `nice` command assigns a priority to a process when it starts.
- The `renice` command modifies the priority of a running process.

Chapter 6

Manage SELinux Security

Goal

Protect and manage server security by using SELinux.

Objectives

- Explain how SELinux protects resources, change the current SELinux mode of a system, and set the default SELinux mode of a system.
- Manage the SELinux policy rules that determine the default context for files and directories with the `semanage fcontext` command, and apply the context defined by the SELinux policy to files and directories with the `restorecon` command.
- Activate and deactivate SELinux policy rules with the `setsebool` command, manage the persistent value of SELinux Booleans with the `semanage boolean -l` command, and consult `man` pages that end with `_selinux` to find useful information about SELinux Booleans.
- Use SELinux log analysis tools and display useful information during SELinux troubleshooting with the `sealert` command.

Sections

- Change the SELinux Enforcement Mode (and Guided Exercise)
- Control SELinux File Contexts (and Guided Exercise)
- Adjust SELinux Policy with Booleans (and Guided Exercise)
- Investigate and Resolve SELinux Issues (and Guided Exercise)

Lab

- Manage SELinux Security

Change the SELinux Enforcement Mode

Objectives

- Explain how SELinux protects resources, change the current SELinux mode of a system, and set the default SELinux mode of a system.

SELinux Architecture

Security Enhanced Linux (SELinux) is a critical security feature of Linux. Access to files, ports, and other resources is controlled at a granular level. Processes are permitted to access only the resources that their SELinux policy or Boolean settings specify.

File permissions control file access for a specific user or group. However, file permissions do not prevent an authorized user with file access from using a file for an unintended purpose.

For example, with write access to a file, other editors or programs can still open and modify a structured data file that is designed for only a specific program to write to, which could result in corruption or a data security issue. File permissions do not stop such undesired access, because they do not control *how* a file is used but only *who* is allowed to read, write, or run a file.

SELinux consists of application-specific policies that the application's developers define to declare what actions and accesses are allowed for each binary executable, configuration file, and data file that the application uses. This policy is known as a *targeted policy*, because one policy defines an application's activities. Policies declare the predefined labels that are configured on individual programs, files, and network ports.

SELinux Usage

SELinux enforces a set of access rules that explicitly define allowed actions between processes and resources. Any action that is not defined in an access rule is not allowed. Because only defined actions are allowed, applications with a poor security design are still protected from malicious use. Applications or services with a targeted policy run in a *confined* domain, whereas an application without a policy runs *unconfined* but without any SELinux protection. Individual targeted policies can be disabled to assist with application and security policy development and debugging.

SELinux has the following operational modes:

- **Enforcing** : SELinux enforces the loaded policies. This mode is the default in Red Hat Enterprise Linux.
- **Permissive** : SELinux loads the policies and is active, but instead of enforcing access control rules, it logs access violations. This mode is helpful for testing and troubleshooting applications and rules.
- **Disabled** : SELinux is turned off. SELinux violations are not denied or logged. Disabling SELinux is strongly discouraged.



Important

Starting in Red Hat Enterprise Linux 9, SELinux can be fully disabled only by using the `selinux=0` kernel parameter at boot. RHEL no longer supports setting the `SELINUX=disabled` option in the `/etc/selinux/config` file.

Starting in RHEL 9, disabling SELinux in the `/etc/selinux/config` file results in SELinux starting and performing active enforcement, but without loading any policies. Because policy rules define allowed actions, if no policies are loaded then all actions are denied. This behavior is intentional, and is designed to block malicious attempts to circumvent SELinux protection.

Basic SELinux Concepts

The primary goal of SELinux is to protect user data from improper use by compromised applications or system services. Most Linux administrators are familiar with the standard user, group, and world file permission security model, which is known as *Discretionary Access Control* (DAC) because administrators set file permissions as they need. SELinux provides an additional layer of object-based security, which is defined in granular rules, which are known as *Mandatory Access Control* (MAC) because MAC policies apply to all users and cannot be bypassed for specific users by discretionary configuration settings.

For example, a web server's open firewall port allows remote anonymous access to a web client. However, a malicious user that accesses that port might try to compromise a system through an existing vulnerability. If an example vulnerability compromises the permissions for the `apache` user and group, then a malicious user might directly access the `/var/www/html` document root content, or the system's `/tmp` and `/var/tmp` directories, or other accessible files and directories.

SELinux policies are security rules that define how specific processes access relevant files, directories, and ports. Every resource entity, such as a file, process, directory, or port, has a label called an *SELinux context*. The context label matches a defined SELinux policy rule to allow a process to access the labeled resource. By default, an SELinux policy does not allow any access unless an explicit rule grants access. When no allow rule is defined, all access is disallowed.

SELinux labels have `user`, `role`, `type`, and `security level` fields. Targeted policy, which is enabled in RHEL by default, defines rules by using the `type` context. Type context names typically end with `_t`.



Figure 6.1: SELinux file context

Policy Access Rule Concepts

For example, a web server process is labeled with the `httpd_t` type context. Web server files and directories in the `/var/www/html/` directory and other locations are labeled with the `httpd_sys_content_t` type context. Temporary files in the `/tmp` and `/var/tmp` directories have the `tmp_t` type contexts as a label. The web server's ports have the `http_port_t` type context as a label.

An Apache web server process runs with the `httpd_t` type context. A policy rule permits the Apache server to access files and directories that are labeled with the `httpd_sys_content_t` type context. By default, files in the `/var/www/html` directory have

the `httpd_sys_content_t` type context. A web server policy has by default no allow rules for using files that are labeled `tmp_t`, such as in the `/tmp` and `/var/tmp` directories, thus disallowing access. With SELinux enabled, a malicious user who uses a compromised Apache process would still not have access to the `/tmp` directory files.

A MariaDB server process runs with the `mysqld_t` type context. By default, files in the `/data/mysql` directory have the `mysqld_db_t` type context. A MariaDB server can access the `mysqld_db_t` labeled files, but has no rules to allow access to files for other services, such as `httpd_sys_content_t` labeled files.




Figure 6.2: SELinux decision-making flow

Many commands that list resources use the `-Z` option to manage SELinux contexts. For example, the `ps`, `ls`, `cp`, and `mkdir` commands all use the `-Z` option.

```
[root@host ~]# ps axZ
LABEL                                PID TTY          STAT TIME COMMAND
system_u:system_r:kernel_t:s0        2 ?            S    0:00 [kthreadd]
system_u:system_r:kernel_t:s0        3 ?            I<   0:00 [rcu_gp]
system_u:system_r:kernel_t:s0        4 ?            I<   0:00 [rcu_par_gp]
...output omitted...
[root@host ~]# systemctl start httpd
[root@host ~]# ps -ZC httpd
LABEL                                PID TTY          TIME CMD
system_u:system_r:httpd_t:s0         1550 ?            00:00:00 httpd
system_u:system_r:httpd_t:s0         1551 ?            00:00:00 httpd
system_u:system_r:httpd_t:s0         1552 ?            00:00:00 httpd
system_u:system_r:httpd_t:s0         1553 ?            00:00:00 httpd
system_u:system_r:httpd_t:s0         1554 ?            00:00:00 httpd
```

```
[root@host ~]# ls -Z /var/www
system_u:object_r:httpd_sys_script_exec_t:s0 cgi-bin
system_u:object_r:httpd_sys_content_t:s0 html
```

Change the SELinux Mode

Use the `getenforce` command to view the current SELinux mode. Use the `setenforce` command to change the SELinux mode.

```
[root@host ~]# getenforce
Enforcing
[root@host ~]# setenforce
usage: setenforce [ Enforcing | Permissive | 1 | 0 ]
[root@host ~]# setenforce 0
[root@host ~]# getenforce
Permissive
[root@host ~]# setenforce Enforcing
[root@host ~]# getenforce
Enforcing
```

Alternatively, set the SELinux mode at boot time with a kernel parameter. Pass the `enforcing=0` kernel parameter to boot the system into `permissive` mode, or pass `enforcing=1` to boot into `enforcing` mode. Disable SELinux by passing the `selinux=0` kernel parameter, or pass `selinux=1` to enable SELinux.

Red Hat recommends rebooting the server when you change the SELinux mode from `Permissive` to `Enforcing`. This reboot ensures that the services that are started in `permissive` mode are confined in the next boot.

Set the Default SELinux Mode

To configure SELinux persistently, use the `/etc/selinux/config` file. In the following default example, the configuration sets SELinux to the `enforcing` mode. The comments list other valid values, such as the `permissive` and `disabled` modes.

```
# 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.
...output omitted...
#
# NOTE: In earlier Fedora kernel builds, SELINUX=disabled would also
# fully disable SELinux during boot. If you need a system with SELinux
# fully disabled instead of SELinux running with no policy loaded, you
# need to pass selinux=0 to the kernel command line. You can use grubby
# to persistently set the bootloader to boot with selinux=0:
#
#   grubby --update-kernel ALL --args selinux=0
#
# To revert back to SELinux enabled:
#
#   grubby --update-kernel ALL --remove-args selinux
#
```


SELINUX=enforcing

```
# SELINUXTYPE= can take one of these three values:
#   targeted - Targeted processes are protected,
#   minimum - Modification of targeted policy. Only selected processes are
#             protected.
#   mls - Multi Level Security protection.
SELINUXTYPE=targeted
```

The system reads this file at boot time and starts SELinux accordingly. The `selinux=0|1` and `enforcing=0|1` kernel arguments override this configuration.



References

`getenforce(8)`, `setenforce(8)`, and `selinux_config(5)` man pages

► Guided Exercise

Change the SELinux Enforcement Mode

Manage SELinux modes, both temporarily and persistently.

Outcomes

- View and set the current SELinux mode.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start selinux-opsmode
```

Instructions

- 1. On the `workstation` machine, use the `ssh` command to log in to the `servera` machine as the `student` user and then switch to the `root` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

- 2. Change the default SELinux mode to permissive.

- 2.1. Use the `getenforce` command to verify the current SELinux mode on the `servera` machine.

```
[root@servera ~]# getenforce
Enforcing
```

- 2.2. Use the `vim /etc/selinux/config` command to edit the configuration file. Change the `SELINUX` parameter from `enforcing` to `permissive` mode.

```
[root@servera ~]# vim /etc/selinux/config
```

- 2.3. Use the `grep` command to confirm that the `SELINUX` parameter displays the `permissive` mode.

```
[root@servera ~]# grep '^SELINUX' /etc/selinux/config
SELINUX=permissive
SELINUXTYPE=targeted
```

- 2.4. Use the `setenforce` command to change the SELINUX parameter to the `permissive` mode and verify the change.

```
[root@servera ~]# setenforce 0
[root@servera ~]# getenforce
Permissive
```

- ▶ 3. Change the default SELinux mode back to the `enforcing` mode in the configuration file.

- 3.1. Use the `vim /etc/selinux/config` command to edit the configuration file. Change the SELINUX parameter from `permissive` to `enforcing` mode.

```
[root@servera ~]# vim /etc/selinux/config
```

- 3.2. Use the `grep` command to confirm that the SELINUX parameter sets the `enforcing` mode on booting.

```
[root@servera ~]# grep '^SELINUX' /etc/selinux/config
SELINUX=enforcing
SELINUXTYPE=targeted
```

- ▶ 4. Set the SELinux mode to `enforcing` on the command line. Reboot the `servera` machine and verify the SELinux mode.

- 4.1. Use the `setenforce` command to set the current SELinux mode to the `enforcing` mode. Use the `getenforce` command to confirm that SELinux is set to the `enforcing` mode.

```
[root@servera ~]# setenforce 1
[root@servera ~]# getenforce
Enforcing
```

- 4.2. Reboot the `servera` machine to implement the persistent configuration.

```
[root@servera ~]# systemctl reboot
Connection to servera closed by remote host.
Connection to servera closed.
[student@workstation ~]$
```

- 4.3. Log in to the `servera` machine and verify the SELinux mode.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]# getenforce
Enforcing
```

- 5. Return to the workstation machine as the student user.

```
[root@servera ~]# exit
logout
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the workstation machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish selinux-opsmode
```

Control SELinux File Contexts

Objectives

- Manage the SELinux policy rules that determine the default context for files and directories with the `semanage fcontext` command and apply the context defined by the SELinux policy to files and directories with the `restorecon` command.

Initial SELinux Context

All resources, such as processes, files, and ports, are labeled with an SELinux *context*. SELinux maintains a file-based database of file labeling policies in the `/etc/selinux/targeted/contexts/files/` directory. New files obtain a default label when their file name matches an existing labeling policy.

When a new file's name does not match an existing labeling policy, the file inherits the same label as the parent directory. With labeling inheritance, all files are always labeled when created, regardless of whether an explicit policy exists for a file.

When files are created in default locations that have an existing labeling policy, or when a policy exists for a custom location, then new files are labeled with a correct SELinux context. However, if a file is created in an unexpected location without an existing labeling policy, then the inherited label might not be correct for the new file's intended purpose.

Furthermore, copying a file to a new location can cause that file's SELinux context to change, where the new context is determined by the new location's labeling policy, or from parent directory inheritance if no policy exists. A file's SELinux context can be preserved during copying to retain the context label that was determined for the file's original location. For example, the `cp -p` command preserves all file attributes where possible, and the `cp --preserve=context` command preserves only SELinux contexts, during copying.



Note

Copying a file always creates a file inode, and that inode's attributes, including the SELinux context, must be initially set, as previously discussed.

However, moving a file does not typically create an inode if the move occurs within the same file system, but instead moves the existing inode's file name to a new location. Because the existing inode's attributes do not need to be initialized, a file that is moved with `mv` preserves its SELinux context unless you set a new context on the file with the `-Z` option.

After you copy or move a file, verify that it has the appropriate SELinux context and set it correctly if necessary.

The following example demonstrates how this process works.

Create two files in the `/tmp` directory. Both files receive the `user_tmp_t` context type.

Move the first file, and copy the second file, to the `/var/www/html` directory.

- The moved file retains the file context that was labeled from the original `/tmp` directory.
- The copied file has a new inode and inherits the SELinux context from the destination `/var/www/html` directory.

The `ls -Z` command displays the SELinux context of a file. Observe the label of the files that are created in the `/tmp` directory.

```
[root@host ~]# touch /tmp/file1 /tmp/file2
[root@host ~]# ls -Z /tmp/file*
unconfined_u:object_r:user_tmp_t:s0 /tmp/file1
unconfined_u:object_r:user_tmp_t:s0 /tmp/file2
```

The `ls -Zd` command displays the SELinux context of the specified directory. Note the label on the `/var/www/html` directory and the files inside it.

```
[root@host ~]# ls -Zd /var/www/html/
system_u:object_r:httpd_sys_content_t:s0 /var/www/html/
[root@host ~]# ls -Z /var/www/html/index.html
unconfined_u:object_r:httpd_sys_content_t:s0 /var/www/html/index.html
```

Move one file from the `/tmp` directory to the `/var/www/html` directory. Copy the other file to the same directory. Note the resulting label on each file.

```
[root@host ~]# mv /tmp/file1 /var/www/html/
[root@host ~]# cp /tmp/file2 /var/www/html/
[root@host ~]# ls -Z /var/www/html/file*
unconfined_u:object_r:user_tmp_t:s0 /var/www/html/file1
unconfined_u:object_r:httpd_sys_content_t:s0 /var/www/html/file2
```

The moved file retained its original label and the copied file inherited the destination directory label. The `unconfined_u` is the SELinux user role, `object_r` is the SELinux role, and `s0` is the (lowest possible) sensitivity level. Advanced SELinux configurations and features use these values.

Change the SELinux Context

You can manage the SELinux context on files with the `semanage fcontext`, `restorecon`, and `chcon` commands.

The recommended method to change the context for a file is to create a file context policy by using the `semanage fcontext` command, and then to apply the specified context in the policy to the file by using the `restorecon` command. This method ensures that you can relabel the file to its correct context with the `restorecon` command whenever necessary. The advantage of this method is that you do not need to remember what the context is supposed to be, and you can correct the context on a set of files.

The `chcon` command changes the SELinux context directly on files, without referencing the system's SELinux policy. Although `chcon` is useful for testing and debugging, changing contexts manually with this method is temporary. Although file contexts that you can change manually survive a reboot, they might be replaced if you run `restorecon` to relabel the contents of the file system.

**Important**

When an SELinux system *relabel* occurs, all files on a system are labeled with their policy defaults. When you use `restorecon` on a file, any context that you change manually on the file is replaced if it does not match the rules in the SELinux policy.

The following example creates a directory with a `default_t` SELinux context, which it inherited from the `/` parent directory.

```
[root@host ~]# mkdir /virtual
[root@host ~]# ls -Zd /virtual
unconfined_u:object_r:default_t:s0 /virtual
```

The `chcon` command sets the file context of the `/virtual` directory to the `httpd_sys_content_t` type.

```
[root@host ~]# chcon -t httpd_sys_content_t /virtual
[root@host ~]# ls -Zd /virtual
unconfined_u:object_r:httpd_sys_content_t:s0 /virtual
```

Running the `restorecon` command resets the context to the default value of `default_t`. Note the `Relabeled` message.

```
[root@host ~]# restorecon -v /virtual
Relabeled /virtual from unconfined_u:object_r:httpd_sys_content_t:s0 to
unconfined_u:object_r:default_t:s0
[root@host ~]# ls -Zd /virtual
unconfined_u:object_r:default_t:s0 /virtual
```

Define SELinux Default File Context Policies

The `semanage fcontext` command displays and modifies the policies that determine the default file contexts. You can list all the file context policy rules by running the `semanage fcontext -l` command. These rules use extended regular expression syntax to specify the path and file names.

When viewing policies, the most common extended regular expression is `(/.*)?`, which is usually appended to a directory name. This notation is humorously called *the pirate*, because it looks like a face with an eye patch and a hooked hand next to it.

This syntax is described as "a set of characters that begin with a slash and followed by any number of characters, where the set can either exist or not exist". Stated more simply, this syntax matches the directory itself, even when empty, and also matches almost any file name that is created within that directory.

For example, the following rule specifies that the `/var/www/cgi-bin` directory, and any files in it or in its subdirectories (and in their subdirectories, and so on), have the `system_u:object_r:httpd_sys_script_exec_t:s0` SELinux context, unless a more specific rule overrides this one.

```
/var/www/cgi-bin(/.*)? all files system_u:object_r:httpd_sys_script_exec_t:s0
```



Note

The `all files` field option from the previous example is the default file type that `semanage` uses when you do not specify one. This option applies to all file types that you can use with `semanage`; they are the same as the standard file types as in the *Control Access to Files* chapter in the *Red Hat System Administration I (RH124)* course. You can get more information from the `semanage-fcontext(8)` man page.

Basic File Context Operations

The following table is a reference for the `semanage fcontext` command options to add, remove, or list SELinux file context policies.

The `semanage fcontext` Command

Option	Description
<code>-a, --add</code>	Add a record of the specified object type.
<code>-d, --delete</code>	Delete a record of the specified object type.
<code>-l, --list</code>	List records of the specified object type.

To manage SELinux contexts, install the `policycoreutils` and `policycoreutils-python-utils` packages, which contain the `restorecon` and `semanage` commands.

To reset all files in a directory to the default policy context, first use the `semanage fcontext -l` command to locate and verify that the correct policy exists with the intended file context. Then, use the `restorecon` command on the wildcarded directory name to reset all the files recursively. In the following example, view the file contexts before and after using the `semanage` and `restorecon` commands.

First, verify the SELinux context for the files:

```
[root@host ~]# ls -Z /var/www/html/file*
unconfined_u:object_r:user_tmp_t:s0 /var/www/html/file1
unconfined_u:object_r:htpdp_sys_content_t:s0 /var/www/html/file2
```

Then, use the `semanage fcontext -l` command to list the default SELinux file contexts:

```
[root@host ~]# semanage fcontext -l
...output omitted...
/var/www(/.*)?      all files      system_u:object_r:htpdp_sys_content_t:s0
...output omitted...
```

The `semanage` command output indicates that all the files and subdirectories in the `/var/www/` directory have the `htpdp_sys_content_t` context by default. Running `restorecon` command on the wildcarded directory restores the default context on all files and subdirectories.


```
[root@host ~]# restorecon -Rv /var/www/
Relabeled /var/www/html/file1 from unconfined_u:object_r:user_tmp_t:s0 to
unconfined_u:object_r:httpd_sys_content_t:s0
[root@host ~]# ls -Z /var/www/html/file*
unconfined_u:object_r:httpd_sys_content_t:s0 /var/www/html/file1
unconfined_u:object_r:httpd_sys_content_t:s0 /var/www/html/file2
```

The following example uses the `semanage` command to add a context policy for a new directory. First, create the `/virtual` directory with an `index.html` file inside it. View the SELinux context for the file and the directory.

```
[root@host ~]# mkdir /virtual
[root@host ~]# touch /virtual/index.html
[root@host ~]# ls -Zd /virtual/
unconfined_u:object_r:default_t:s0 /virtual/
[root@host ~]# ls -Z /virtual/
unconfined_u:object_r:default_t:s0 index.html
```

Next, use the `semanage fcontext` command to add an SELinux file context policy for the directory.

```
[root@host ~]# semanage fcontext -a -t httpd_sys_content_t '/virtual(/.*)?'
```

Use the `restorecon` command on the wildcarded directory to set the default context on the directory and all files within it.

```
[root@host ~]# restorecon -RFv /virtual
Relabeled /virtual from unconfined_u:object_r:default_t:s0 to
system_u:object_r:httpd_sys_content_t:s0
Relabeled /virtual/index.html from unconfined_u:object_r:default_t:s0 to
system_u:object_r:httpd_sys_content_t:s0
[root@host ~]# ls -Zd /virtual/
drwxr-xr-x. root root system_u:object_r:httpd_sys_content_t:s0 /virtual/
[root@host ~]# ls -Z /virtual/
-rw-r--r--. root root system_u:object_r:httpd_sys_content_t:s0 index.html
```

Use the `semanage fcontext -l -C` command to view any local customizations to the default policy.

```
[root@host ~]# semanage fcontext -l -C
SELinux fcontext      type      Context

/virtual(/.*)?       all files system_u:object_r:httpd_sys_content_t:s0
```



References

`chcon(1)`, `restorecon(8)`, `semanage(8)`, and `semanage-fcontext(8)` man pages

► Guided Exercise

Control SELinux File Contexts

Persistently change the SELinux context of a directory and its contents.

Outcomes

- Configure the Apache HTTP server to publish web content from a non-standard document root.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start selinux-filecontexts
```

Instructions

- 1. Log in to `servera` as the `student` user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

- 2. Configure Apache to use a document directory in a non-standard location.

- 2.1. Create the `/custom` directory.

```
[root@servera ~]# mkdir /custom
```

- 2.2. Create the `index.html` file in the `/custom` directory that contains the `This is SERVERA.` text.

```
[root@servera ~]# echo 'This is SERVERA.' > /custom/index.html
```

- 2.3. Configure Apache to use the new directory location. Edit the Apache `/etc/httpd/conf/httpd.conf` configuration file, and replace the two occurrences of the `/var/www/html` directory with the `/custom` directory. You can use the `vim /etc/httpd/conf/httpd.conf` command to do so. The following example shows the expected content of the `/etc/httpd/conf/httpd.conf` file.

```
[root@servera ~]# cat /etc/httpd/conf/httpd.conf
...output omitted...
DocumentRoot "/custom"
...output omitted...
<Directory "/custom">
...output omitted...
```

- ▶ 3. Start and enable the Apache web service and confirm that the service is running.

3.1. Start and enable the Apache web service by using the `systemctl` command.

```
[root@servera ~]# systemctl enable --now httpd
Created symlink /etc/systemd/system/multi-user.target.wants/httpd.service → /usr/lib/systemd/system/httpd.service.
```

3.2. Verify that the service is running.

```
[root@servera ~]# 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 Wed 2022-04-06 05:21:19 EDT; 22s ago
  Docs: man:httpd.service(8)
  Main PID: 1676 (httpd)
  ...output omitted...
Apr 06 05:21:19 servera.lab.example.com systemd[1]: Starting The Apache HTTP Server...
Apr 06 05:21:19 servera.lab.example.com systemd[1]: Started The Apache HTTP Server.
Apr 06 05:21:19 servera.lab.example.com httpd[1676]: Server configured, listening on: port 80
```

- ▶ 4. Open a web browser on workstation and try to view the `http://servera/index.html` web page. You get an error message that you do not have permission to access the file.
- ▶ 5. To grant access to the `index.html` file on `servera`, you must configure the SELinux context. Define an SELinux file context rule that sets the context type to `httpd_sys_content_t` for the `/custom` directory and all the files under it.

```
[root@servera ~]# semanage fcontext -a \
-t httpd_sys_content_t '/custom(/.*)?'
```

- ▶ 6. Correct the file contexts in the `/custom` directory.

```
[root@servera ~]# restorecon -Rv /custom
Relabeled /custom from unconfined_u:object_r:default_t:s0 to
unconfined_u:object_r:httpd_sys_content_t:s0
Relabeled /custom/index.html from unconfined_u:object_r:default_t:s0 to
unconfined_u:object_r:httpd_sys_content_t:s0
```

- ▶ 7. Try to view `http://servera/index.html` again in the web browser on the workstation machine. You should see the `This is SERVERA.` message.
- ▶ 8. Return to the workstation machine as the student user.

```
[root@servera ~]# exit
logout
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the workstation machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish selinux-filecontexts
```

Adjust SELinux Policy with Booleans

Objectives

- Activate and deactivate SELinux policy rules with the `setsebool` command, manage the persistent value of SELinux Booleans with the `semanage boolean -l` command, and consult man pages that end with `_selinux` to find useful information about SELinux Booleans.

SELinux Booleans

An application or service developer writes an SELinux targeted policy to define the allowed behavior of the targeted application. A developer can include optional application behavior in the SELinux policy that can be enabled when the behavior is allowed on a specific system. SELinux Booleans enable or disable the SELinux policy's optional behavior. With Booleans, you can selectively tune the behavior of an application.

These optional behaviors are application-specific, and must be discovered and selected for each targeted application. Service-specific Booleans are documented in that service's SELinux man page. For example, the web server `httpd` service has its `httpd(8)` man page, and an `httpd_selinux(8)` man page to document its SELinux policy, including the supported process types, file contexts, and the available Boolean-enabled behaviors. The SELinux man pages are provided in the `selinux-policy-doc` package.

Use the `getsebool` command to list available Booleans for the targeted policies on this system, and the current Boolean status. Use the `setsebool` command to enable or disable the running state of these behaviors. The `setsebool -P` command option makes the setting persistent by writing to the policy file. Only privileged users can set SELinux Booleans.

```
[root@host ~]# getsebool -a
abrt_anon_write --> off
abrt_handle_event --> off
abrt_upload_watch_anon_write --> on
...output omitted...
```

Example httpd Policy Boolean

The `httpd` service policy includes the `httpd_enable_homedirs` Boolean, which enables the sharing of home directories with `httpd`. Typically, a user's local home directory is accessible to the user only when logged in to the local system. Alternatively, home directories are shared and accessed by using a remote file sharing protocol, such as NFS. In both scenarios, home directories are not shared by using `https`, by default, and are not available to the user through a browser.

```
[root@host ~]# getsebool httpd_enable_homedirs
httpd_enable_homedirs --> off
```

You can enable sharing and enable users to access their home directories with a browser. When enabled, the `httpd` service shares home directories that are labeled with the `user_home_dir_t` file context. Users can then access and manage their home directory files from a browser.

Manage the Policy Boolean

Setting SELinux Booleans with the `setsebool` command without the `-P` option is temporary, and settings return to the persistent values after rebooting. View additional information with the `semanage boolean -l` command, which lists the Booleans from the policy files, including whether a Boolean is persistent, the default and current values, and a short description.

```
[root@host ~]# semanage boolean -l | grep httpd_enable_homedirs
httpd_enable_homedirs      (off , off) Allow httpd to enable homedirs
[root@host ~]# setsebool httpd_enable_homedirs on
[root@host ~]# semanage boolean -l | grep httpd_enable_homedirs
httpd_enable_homedirs      (on , off) Allow httpd to enable homedirs
[root@host ~]# getsebool httpd_enable_homedirs
httpd_enable_homedirs --> on
```

To list only Booleans with a current setting that is different from the default setting at boot, use the `semanage boolean -l -C` command. This example has the same result as the previous example, without requiring the `grep` filtering.

```
[root@host ~]# semanage boolean -l -C
SELinux boolean           State Default Description

httpd_enable_homedirs     (on , off) Allow httpd to enable homedirs
```

The previous example temporarily set the current value for the `httpd_enable_homedirs` Boolean to `on`, until the system reboots. To change the default setting, use the `setsebool -P` command to make the setting persistent. The following example sets a persistent value, and then views the Boolean's information from the policy file.

```
[root@host ~]# setsebool -P httpd_enable_homedirs on
[root@host ~]# semanage boolean -l | grep httpd_enable_homedirs
httpd_enable_homedirs     (on , on) Allow httpd to enable homedirs
```

Use the `semanage boolean -l -C` command again. The Boolean is displayed despite the appearance that the current and default settings are the same. However, the `-C` option matches when the current setting is different from the default setting from the last boot. For this `httpd_enable_homedirs` example, the original default boot setting was `off`.

```
[root@host ~]# semanage boolean -l -C
SELinux boolean           State Default Description

httpd_enable_homedirs     (on , on) Allow httpd to enable homedirs
```



References

`booleans(8)`, `getsebool(8)`, `setsebool(8)`, `semanage(8)`, and `semanage-boolean(8)` man pages

► Guided Exercise

Adjust SELinux Policy with Booleans

Configure Apache to publish web content from users' home directories.

Outcomes

- Configure the Apache web service to publish web content from the user's home directory.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start selinux-booleans
```

Instructions

- 1. On the `workstation` machine, use the `ssh` command to log in to the `servera` machine as the `student` user and then switch to the `root` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

- 2. Edit the `/etc/httpd/conf.d/userdir.conf` configuration file to enable the Apache feature so that users can publish web content from their home directory. Comment out the line in the `IfModule` section that sets the `UserDir` variable to the `disabled` value, and uncomment the line that sets the `UserDir` variable to the `public_html` value.

```
[root@servera ~]# vim /etc/httpd/conf.d/userdir.conf
<IfModule mod_userdir.c>
...output omitted...
    # UserDir disabled

...output omitted...
    UserDir public_html

...output omitted...
</IfModule>
```

- 3. Start and enable the Apache web service.

```
[root@servera ~]# systemctl enable --now httpd
```

- ▶ 4. Open another terminal window, and use the `ssh` command to log in to the `servera` machine as the `student` user. Create the `index.html` web content file in the `~/public_html` directory.

- 4.1. In another terminal window, use the `ssh` command to log in to the `servera` machine as the `student` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
```

- 4.2. Use the `mkdir` command to create the `~/public_html` directory.

```
[student@servera ~]$ mkdir ~/public_html
```

- 4.3. Create the `index.html` file with the following content:

```
[student@servera ~]$ echo 'This is student content on SERVERA.' > \
~/public_html/index.html
```

- 4.4. For the Apache web service to serve the contents of the `/home/student/public_html` directory, it must be allowed to share files and subdirectories in the `/home/student` directory. When you created the `/home/student/public_html` directory, it was automatically configured to allow anyone with home directory permission to access its contents.

Change the `/home/student` directory permissions to allow the Apache web service to access the `public_html` subdirectory.

```
[student@servera ~]$ chmod 711 ~
[student@servera ~]$ ls -ld ~
drwx--x--x. 16 student student 4096 Nov  3 09:28 /home/student
```

- ▶ 5. Open a web browser on the `workstation` machine and enter the `http://servera/~student/index.html` address. An error message states that you do not have permission to access the file.
- ▶ 6. Switch to the other terminal, and use the `getsebool` command to see whether any Booleans restrict access to home directories for the `httpd` service.

```
[root@servera ~]# getsebool -a | grep home
...output omitted...
httpd_enable_homedirs --> off
...output omitted...
```

- ▶ 7. Use the `setsebool` command to enable persistent access to the home directory for the `httpd` service.


```
[root@servera ~]# setsebool -P httpd_enable_homedirs on
```

- ▶ 8. Verify that you can now see the `This is student content on SERVERA.` message in the web browser after entering the `http://servera/~student/index.html` address. You might need to close and reopen your web browser to see the message.
- ▶ 9. Return to the `workstation` machine as the `student` user.

```
[root@servera ~]# exit
logout
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish selinux-boolean
```

Investigate and Resolve SELinux Issues

Objectives

- Use SELinux log analysis tools and display useful information during SELinux troubleshooting with the `sealert` command.

Troubleshoot SELinux Issues

When applications unexpectedly fail to work due to SELinux access denials, methods and tools are available to resolve these issues. It is helpful to start by understanding some fundamental concepts and behaviors when SELinux is enabled.

- SELinux consists of targeted policies that explicitly define allowable actions.
- A policy entry defines a labeled process and a labeled resource that interact.
- The policy states the process type, and the file or port context, by using labels.
- The policy entry defines one process type, one resource label, and the explicit action to allow.
- An action can be a system call, a kernel function, or another specific programming routine.
- If no entry is created for a specific process-resource-action relationship, then the action is denied.
- When an action is denied, the attempt is logged with useful context information.

Red Hat Enterprise Linux provides a stable targeted SELinux policy for almost every service in the distribution. Therefore, it is unusual to have SELinux access problems with common RHEL services when they are configured correctly. SELinux access problems occur when services are implemented incorrectly, or when new applications have incomplete policies. Consider these troubleshooting concepts before making broad SELinux configuration changes.

- Most access denials indicate that SELinux is working correctly by blocking improper actions.
- Evaluating denied actions requires some familiarity with normal, expected service actions.
- The most common SELinux issue is an incorrect context on new, copied, or moved files.
- File contexts can be fixed when an existing policy references their location.
- Optional Boolean policy features are documented in the `_selinux` man pages.
- Implementing Boolean features generally requires setting additional non-SELinux configuration.
- SELinux policies do not replace or circumvent file permissions or access control list restrictions.

When a common application or service fails, and the service is known to have a working SELinux policy, first see the service's `_selinux` man page to verify the correct context type label. View the affected process and file attributes to verify that the correct labels are set.

Monitor SELinux Violations

The SELinux troubleshooting service, from the `setroubleshoot-server` package, provides tools to diagnose SELinux issues. When SELinux denies an action, an Access Vector Cache (AVC) message is logged to the `/var/log/audit/audit.log` security log file. The SELinux troubleshooting service monitors for AVC events and sends an event summary to the `/var/log/messages` file.

The AVC summary includes an event unique identifier (UUID). Use the `sealert -l UUID` command to view comprehensive report details for the specific event. Use the `sealert -a /var/log/audit/audit.log` command to view all existing events.

Consider the following example sequence of commands on a standard Apache web server. You create `/root/mypage` and move it to the default Apache content directory (`/var/www/html`). Then, after starting the Apache service, you try to retrieve the file content.

```
[root@host ~]# touch /root/mypage
[root@host ~]# mv /root/mypage /var/www/html
[root@host ~]# systemctl start httpd
[root@host ~]# curl http://localhost/mypage
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
<html><head>
<title>403 Forbidden</title>
</head><body>
<h1>Forbidden</h1>
<p>You don't have permission to access this resource.</p>
</body></html>
```

The web server does not display the content, and returns a `permission denied` error. An AVC event is logged to the `/var/log/audit/audit.log` and `/var/log/messages` files. Note the suggested `sealert` command and UUID in the `/var/log/messages` event message.

```
[root@host ~]# tail /var/log/audit/audit.log
...output omitted...
type=AVC msg=audit(1649249057.067:212): avc: denied { getattr }
for pid=2332 comm="httpd" path="/var/www/html/mypage"
dev="vda4" ino=9322502 scontext=system_u:system_r:httpd_t:s0
tcontext=unconfined_u:object_r:admin_home_t:s0 tclass=file permissive=0
...output omitted
[root@host ~]# tail /var/log/messages
...output omitted...
Apr  6 08:44:19 host setroubleshoot[2547]: SELinux is preventing /usr/sbin/httpd
from getattr access on the file /var/www/html/mypage. For complete SELinux
messages run: sealert -l 95f41f98-6b56-45bc-95da-ce67ec9a9ab7
...output omitted...
```

The `sealert` output describes the event, and includes the affected process, the accessed file, and the attempted and denied action. The output includes advice for correcting the file's label, if appropriate. Additional advice describes how to generate a new policy to allow the denied action. Use the given advice only when it is appropriate for your scenario.



Important

The `sealert` output includes a confidence rating, which indicates the level of confidence that the given advice will mitigate the denial. However, that advice might not be appropriate for your scenario.

For example, if the AVC denial is because the denied file is in the wrong location, then advice that states either to adjust the file's context label, or to create a policy for this location and action, although technically accurate, is not the correct solution for your scenario. If the root cause is a wrong location or file name, then moving or renaming the file and then restoring a correct file context is the correct solution instead.

```
[root@host ~]# sealert -l 95f41f98-6b56-45bc-95da-ce67ec9a9ab7
SELinux is preventing /usr/sbin/httpd from getattr access on the file /var/www/html/mypage.

**** Plugin restorecon (99.5 confidence) suggests ****

If you want to fix the label.
/var/www/html/mypage default label should be httpd_sys_content_t.
Then you can run restorecon. The access attempt may have been stopped due to
insufficient permissions to access a parent directory in which case try to change
the following command accordingly.
Do
# /sbin/restorecon -v /var/www/html/mypage

**** Plugin catchall (1.49 confidence) suggests ****

If you believe that httpd should be allowed getattr access on the mypage 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:
# ausearch -c 'httpd' --raw | audit2allow -M my-httpd
# semodule -X 300 -i my-httpd.pp

Additional Information:
Source Context                system_u:system_r:httpd_t:s0
Target Context                unconfined_u:object_r:admin_home_t:s0
Target Objects                /var/www/html/mypage [ file ]
Source                        httpd
Source Path                    /usr/sbin/httpd
...output omitted...

Raw Audit Messages
type=AVC msg=audit(1649249057.67:212): avc: denied { getattr }
for pid=2332 comm="httpd" path="/var/www/html/mypage"
dev="vda4" ino=9322502 scontext=system_u:system_r:httpd_t:s0
tcontext=unconfined_u:object_r:admin_home_t:s0 tclass=file permissive=0

type=SYSCALL msg=audit(1649249057.67:212): arch=x86_64 syscall=newfstatat
success=no exit=EACCES a0=ffffff9c a1=7fe9c00048f8 a2=7fe9ccfc8830 a3=100
items=0 ppid=2329 pid=2332 auid=4294967295 uid=48 gid=48 euid=48 suid=48 fsuid=48
egid=48 sgid=48 fsgid=48 tty=(none) ses=4294967295 comm=httpd exe=/usr/sbin/httpd
subj=system_u:system_r:httpd_t:s0 key=(null)

Hash: httpd,httpd_t,admin_home_t,file,getattr
```

In this example, the accessed file is in the correct location, but does not have the correct SELinux file context. The Raw Audit Messages section displays information from the `/var/log/audit/audit.log` event entry. Use the `restorecon /var/www/html/mypage` command to set the correct context label. To correct multiple files recursively, use the `restorecon -R` command on the parent directory.

Use the `ausearch` command to search for AVC events in the `/var/log/audit/audit.log` log file. Use the `-m` option to specify the AVC message type and the `-ts` option to provide a time hint, such as `recent`.

```
[root@host ~]# ausearch -m AVC -ts recent
-----
time->Tue Apr 6 13:13:07 2019
type=PROCTITLE msg=audit(1554808387.778:4002):
  proctitle=2F7573722F7362696E2F6874747064002D44464F524547524F554E44
type=SYSCALL msg=audit(1554808387.778:4002): arch=c000003e syscall=49
  success=no exit=-13 a0=3 a1=55620b8c9280 a2=10 a3=7ffed967661c items=0
  ppid=1 pid=9340 auid=4294967295 uid=0 gid=0 euid=0 suid=0 fsuid=0 egid=0
  sgid=0 fsgid=0 tty=(none) ses=4294967295 comm="httpd" exe="/usr/sbin/httpd"
  subj=system_u:system_r:httpd_t:s0 key=(null)
type=AVC msg=audit(1554808387.778:4002): avc: denied { name_bind }
  for pid=9340 comm="httpd" src=82 scontext=system_u:system_r:httpd_t:s0
  tcontext=system_u:object_r:reserved_port_t:s0 tclass=tcp_socket permissive=0
```

Troubleshoot SELinux Issues with the Web Console

The RHEL web console includes tools for troubleshooting SELinux issues. Select **SELinux** from the menu on the left. The SELinux policy window displays the current enforcing state. The **SELinux access control errors** section lists current SELinux issues.




Figure 6.3: SELinux policy and errors in the web console

Click the `>` character to display event details. Click **solution details** to display all event details and advice. You can click **Apply the solution**.

After correcting the issue, the **SELinux access control errors** section should remove that event from view. If the **No SELinux alerts** message appears, then you have corrected all current SELinux issues.



References

sealert(8) man page

► Guided Exercise

Investigate and Resolve SELinux Issues

Learn how to troubleshoot SELinux security denials.

Outcomes

- Gain experience with SELinux troubleshooting tools.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start selinux-issues
```

Instructions

1. From a web browser on the `workstation` machine, open the `http://servera/index.html` web page. An error message states that you do not have permission to access the file.
2. Use the `ssh` command to log in to `servera` as the `student` user. Use the `sudo -i` command to switch to the `root` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

3. Use the `less` command to view the contents of the `/var/log/messages` file. You use the `/` character and search for the `sealert` text. Press the `n` key until you reach the last occurrence, because previous exercises might also have generated SELinux messages. Copy the suggested `sealert` command so that you can use it in the next step. Use the `q` key to quit the `less` command.

```
[root@servera ~]# less /var/log/messages
...output omitted...
Apr  7 04:52:18 servera setroubleshoot[20715]: SELinux is preventing /usr/sbin/
httpd from getattr access on the file /custom/index.html. For complete SELinux
messages run: sealert -l 9a96294a-239b-4568-8f1e-9f35b5fb472b
...output omitted...
```

- ▶ 4. Run the suggested `sealert` command. Note the source context, the target objects, the policy, and the enforcing mode. Find the correct SELinux context label for the file that the `httpd` service tries to serve.

4.1. Run the `sealert` command.

The output explains that the `/custom/index.html` file has an incorrect context label.

```
[root@servera ~]# sealert -l 9a96294a-239b-4568-8f1e-9f35b5fb472b
SELinux is preventing /usr/sbin/httpd from getattr access on the file /custom/index.html.

**** Plugin catchall_labels (83.8 confidence) suggests ****

If you want to allow httpd to have getattr access on the index.html file
Then you need to change the label on /custom/index.html
Do
# semanage fcontext -a -t FILE_TYPE '/custom/index.html'
where FILE_TYPE is one of the following: NetworkManager_exec_t,
NetworkManager_log_t, NetworkManager_tmp_t, abrt_dump_oops_exec_t,
abrt_etc_t, abrt_exec_t, abrt_handle_event_exec_t, abrt_helper_exec_t,
abrt_retrace_coredump_exec_t, abrt_retrace_spool_t, abrt_retrace_worker_exec_t,
abrt_tmp_t, abrt_upload_watch_tmp_t, abrt_var_cache_t, abrt_var_log_t,
abrt_var_run_t, accountsd_exec_t, acct_data_t, acct_exec_t, admin_crontab_tmp_t,
admin_passwd_exec_t, afs_logfile_t, aide_exec_t, aide_log_t, alsa_exec_t,
alsa_tmp_t, amanda_exec_t, amanda_log_t, amanda_recover_exec_t, amanda_tmp_t,
amtu_exec_t, anacron_exec_t, anon_inodefs_t
...output omitted...

Additional Information:
Source Context      system_u:system_r:httpd_t:s0
Target Context      unconfined_u:object_r:default_t:s0
Target Objects      /custom/index.html [ file ]
Source              httpd
Source Path          /usr/sbin/httpd
Port                <Unknown>
Host                servera.lab.example.com
Source RPM Packages httpd-2.4.51-7.el9_0.x86_64
Target RPM Packages
SELinux Policy RPM  selinux-policy-targeted-34.1.27-1.el9.noarch
Local Policy RPM    selinux-policy-targeted-34.1.27-1.el9.noarch
Selinux Enabled     True
Policy Type         targeted
Enforcing Mode      Enforcing
Host Name           servera.lab.example.com
Platform            Linux servera.lab.example.com
                    5.14.0-70.2.1.el9_0.x86_64 #1 SMP PREEMPT Wed Mar
                    16 18:15:38 EDT 2022 x86_64 x86_64

Alert Count         4
First Seen          2022-04-07 04:51:38 EDT
Last Seen           2022-04-07 04:52:13 EDT
Local ID            9a96294a-239b-4568-8f1e-9f35b5fb472b

Raw Audit Messages
```



```
type=AVC msg=audit(1649321533.406:1024): avc: denied { getattr } for
pid=20464 comm="httpd" path="/custom/index.html" dev="vda4" ino=25571802
scontext=system_u:system_r:httpd_t:s0 tcontext=unconfined_u:object_r:default_t:s0
tclass=file permissive=0
```

...output omitted...

- 4.2. Verify the SELinux context for the directory from where the `httpd` service serves the content by default, `/var/www/html`. The `httpd_sys_content_t` SELinux context is appropriate for the `/custom/index.html` file.

```
[root@servera ~]# ls -ldZ /var/www/html
drwxr-xr-x. 2 root root system_u:object_r:httpd_sys_content_t:s0 6 Mar 21 11:47 /
var/www/html
```

- ▶ 5. The Raw Audit Messages section of the `sealert` command contains information from the `/var/log/audit/audit.log` file. Use the `ausearch` command to search the `/var/log/audit/audit.log` file. The `-m` option searches on the message type. The `-ts` option searches based on time. The following entry identifies the relevant process and file that cause the alert. The process is the `httpd` Apache web server, the file is `/custom/index.html`, and the context is `system_r:httpd_t`.

```
[root@servera ~]# ausearch -m AVC -ts today
...output omitted...
----
time->Thu Apr  7 04:52:13 2022
type=PROCTITLE msg=audit(1649321533.406:1024):
  proctitle=2F7573722F7362696E2F687474064002D44464F524547524F554E44
type=SYSCALL msg=audit(1649321533.406:1024): arch=c000003e syscall=262 success=no
  exit=-13 a0=ffffff9c a1=7fefc403d850 a2=7fefc89bc830 a3=100 items=0 ppid=20461
  pid=20464 auid=4294967295 uid=48 gid=48 euid=48 suid=48 fsuid=48 egid=48
  sgid=48 fsgid=48 tty=(none) ses=4294967295 comm="httpd" exe="/usr/sbin/httpd"
  subj=system_u:system_r:httpd_t:s0 key=(null)
type=AVC msg=audit(1649321533.406:1024): avc: denied
{ getattr } for pid=20464 comm="httpd" path="/custom/index.html"
dev="vda4" ino=25571802 sccontext=system_u:system_r:httpd_t:s0
tcontext=unconfined_u:object_r:default_t:s0 tclass=file permissive=0
```

- ▶ 6. Resolve the issue by applying the `httpd_sys_content_t` context.

```
[root@servera ~]# semanage fcontext -a \
-t httpd_sys_content_t '/custom(/.*)?'
[root@servera ~]# restorecon -Rv /custom
Relabeled /custom from unconfined_u:object_r:default_t:s0 to
unconfined_u:object_r:httpd_sys_content_t:s0
Relabeled /custom/index.html from unconfined_u:object_r:default_t:s0 to
unconfined_u:object_r:httpd_sys_content_t:s0
```

- ▶ 7. Again, try to view `http://servera/index.html`. The `This is SERVERA.` message is displayed.
- ▶ 8. Return to the workstation machine as the student user.

```
[root@servera ~]# exit
logout
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish selinux-issues
```

► Lab

Manage SELinux Security

Identify issues in system log files and adjust the SELinux configuration.

Outcomes

- Identify issues in system log files.
- Adjust the SELinux configuration.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start selinux-review
```

Instructions

1. Log in to the `serverb` machine as the `student` user and switch to the `root` user.
2. From a web browser on the `workstation` machine, view the `http://serverb/lab.html` web page. You see the error message: `You do not have permission to access this resource.`
3. Research and identify the SELinux issue that prevents the Apache service from serving web content.
4. Display the SELinux context of the new HTTP document directory and the original HTTP document directory. Resolve the SELinux issue that prevents the Apache server from serving web content.
5. Verify that the Apache server can now serve web content.
6. Return to the `workstation` machine as the `student` user.

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade selinux-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish selinux-review
```

► Solution

Manage SELinux Security

Identify issues in system log files and adjust the SELinux configuration.

Outcomes

- Identify issues in system log files.
- Adjust the SELinux configuration.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start selinux-review
```

Instructions

1. Log in to the `serverb` machine as the `student` user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
[student@serverb ~]$ sudo -i
[sudo] password for student: student
[root@serverb ~]#
```

2. From a web browser on the workstation machine, view the `http://serverb/lab.html` web page. You see the error message: `You do not have permission to access this resource.`
3. Research and identify the SELinux issue that prevents the Apache service from serving web content.
 - 3.1. View the contents of the `/var/log/messages` file. Use the `/` key and search for the `sealert` string. Use the `q` key to quit the `less` command.

```
[root@serverb ~]# less /var/log/messages
...output omitted...
Apr  7 06:16:15 serverb setroubleshoot[26509]: failed to retrieve rpm info for /
lab-content/la
b.html
Apr  7 06:16:17 serverb setroubleshoot[26509]: SELinux is preventing /usr/sbin/
httpd from getattr access on the file /lab-content/lab.html. For complete SELinux
messages run: sealert -l 35c9e452-2552-4ca3-8217-493b72ba6d0b
Apr  7 06:16:17 serverb setroubleshoot[26509]: SELinux is preventing /usr/sbin/
httpd from getattr access on the file /lab-content/lab.html
...output omitted...
```

3.2. Run the suggested `sealert` command. Note the source context, the target objects, the policy, and the enforcing mode.

```
[root@serverb ~]# sealert -l 35c9e452-2552-4ca3-8217-493b72ba6d0b
SELinux is preventing /usr/sbin/httpd from getattr access on the file /lab-
content/lab.html.

***** Plugin catchall_labels (83.8 confidence) suggests *****

If you want to allow httpd to have getattr access on the lab.html file
Then you need to change the label on /lab-content/lab.html
Do
# semanage fcontext -a -t FILE_TYPE '/lab-content/lab.html'
where FILE_TYPE is one of the following:
...output omitted...

Additional Information:
Source Context      system_u:system_r:httpd_t:s0
Target Context      unconfined_u:object_r:default_t:s0
Target Objects      /lab-content/lab.html [ file ]
Source              httpd
Source Path          /usr/sbin/httpd
Port                <Unknown>
Host                serverb.lab.example.com
Source RPM Packages httpd-2.4.51-7.el9_0.x86_64
Target RPM Packages
SELinux Policy RPM  selinux-policy-targeted-34.1.27-1.el9.noarch
Local Policy RPM    selinux-policy-targeted-34.1.27-1.el9.noarch
Selinux Enabled     True
Policy Type         targeted
Enforcing Mode      Enforcing
Host Name            serverb.lab.example.com
Platform             Linux serverb.lab.example.com
                    5.14.0-70.2.1.el9_0.x86_64 #1 SMP PREEMPT Wed Mar
                    16 18:15:38 EDT 2022 x86_64 x86_64

Alert Count          8
First Seen           2022-04-07 06:14:45 EDT
Last Seen            2022-04-07 06:16:12 EDT
Local ID             35c9e452-2552-4ca3-8217-493b72ba6d0b

Raw Audit Messages
```

```
type=AVC msg=audit(1649326572.86:407): avc: denied { getattr } for
pid=10731 comm="httpd" path="/lab-content/lab.html" dev="vda4" ino=18192752
scontext=system_u:system_r:httpd_t:s0 tcontext=unconfined_u:object_r:default_t:s0
tclass=file permissive=0
```

```
type=SYSCALL msg=audit(1649326572.86:407): arch=x86_64 syscall=newfstatat
success=no exit=EACCES a0=ffffffff a1=7f7c8c0457c0 a2=7f7c887f7830 a3=100 items=0
ppid=10641 pid=10731 auid=4294967295 uid=48 gid=48 euid=48 suid=48 fsuid=48
egid=48 sgid=48 fsgid=48 tty=(none) ses=4294967295 comm=httpd exe=/usr/sbin/httpd
subj=system_u:system_r:httpd_t:s0 key=(null)
```

Hash: httpd,httpd_t,default_t,file,getattr

- 3.3. The Raw Audit Messages section of the `sealert` command contains information from the `/var/log/audit/audit.log` file. Search the `/var/log/audit/audit.log` file. The `-m` option searches on the message type. The `ts` option searches based on time. The following entry identifies the relevant process and file that cause the alert. The process is the `httpd` Apache web server, the file is `/lab-content/lab.html`, and the context is `system_r:httpd_t`.

```
[root@serverb ~]# ausearch -m AVC -ts recent
...output omitted...
----
time->Thu Apr 7 06:16:12 2022
type=PROCTITLE msg=audit(1649326572.086:407):
proctitle=2F7573722F7362696E2F6874747064002D44464F524547524F554E44
type=SYSCALL msg=audit(1649326572.086:407): arch=c000003e syscall=262 success=no
exit=-13 a0=ffffffff a1=7f7c8c0457c0 a2=7f7c887f7830 a3=100 items=0 ppid=10641
pid=10731 auid=4294967295 uid=48 gid=48 euid=48 suid=48 fsuid=48 egid=48
sgid=48 fsgid=48 tty=(none) ses=4294967295 comm="httpd" exe="/usr/sbin/httpd"
subj=system_u:system_r:httpd_t:s0 key=(null)
type=AVC msg=audit(1649326572.086:407): avc: denied { getattr } for
pid=10731 comm="httpd" path="/lab-content/lab.html" dev="vda4" ino=18192752
scontext=system_u:system_r:httpd_t:s0 tcontext=unconfined_u:object_r:default_t:s0
tclass=file permissive=0
```

4. Display the SELinux context of the new HTTP document directory and the original HTTP document directory. Resolve the SELinux issue that prevents the Apache server from serving web content.

- 4.1. Compare the SELinux context for the `/lab-content` and `/var/www/html` directories.

```
[root@serverb ~]# ls -dZ /lab-content /var/www/html
unconfined_u:object_r:default_t:s0 /lab-content
system_u:object_r:httpd_sys_content_t:s0 /var/www/html
```

- 4.2. Create a file context rule that sets the default type to `httpd_sys_content_` for the `/lab-content` directory and all the files in it.

```
[root@serverb ~]# semanage fcontext -a \
-t httpd_sys_content_t '/lab-content(/.*)?'
```

- 4.3. Correct the SELinux context for the files in the `/lab-content` directory.

```
[root@serverb ~]# restorecon -R /lab-content/
```

5. Verify that the Apache server can now serve web content.
 - 5.1. Use your web browser to refresh the `http://serverb/lab.html` link. If the content is displayed, then your issue is resolved.

This is the html file for the SELinux final lab on SERVERB.

6. Return to the workstation machine as the student user.

```
[root@serverb ~]# exit
logout
[student@serverb ~]$ exit
logout
Connection to serverb closed.
[student@workstation ~]$
```

Evaluation

As the student user on the workstation machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade selinux-review
```

Finish

On the workstation machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish selinux-review
```


Summary

- Use the `getenforce` and `setenforce` commands to manage the SELinux mode of a system.
- The `semanage` command manages SELinux policy rules. The `restorecon` command applies the context that the policy defines.
- Booleans are switches that change the behavior of the SELinux policy. You can enable or disable them to tune the policy.
- The `sealert` command displays useful information to help with SELinux troubleshooting.

Chapter 7

Manage Basic Storage

Goal

Create and manage storage devices, partitions, file systems, and swap spaces from the command line.

Objectives

- Create storage partitions, format them with file systems, and mount them for use.
- Create and manage swap spaces to supplement physical memory.

Sections

- Add Partitions, File Systems, and Persistent Mounts (and Guided Exercise)
- Manage Swap Space (and Guided Exercise)

Lab

- Manage Basic Storage

Add Partitions, File Systems, and Persistent Mounts

Objectives

- Create storage partitions, format them with file systems, and mount them for use.

Partition Disks

Disk partitioning divides a hard drive into multiple logical storage *partitions*. You can use partitions to divide storage based on different requirements, and this division provides many benefits:

- Limit available space to applications or users.
- Separate operating system and program files from user files.
- Create a separate area for memory swapping.
- Limit disk space use to improve the performance of diagnostic tools and backup imaging.

MBR Partition Scheme

The *Master Boot Record* (MBR) partitioning scheme is the standard on systems that run BIOS firmware. This scheme supports a maximum of four primary partitions. On Linux systems, with extended and logical partitions, you can create up to 15 partitions. With a 32-bit partition size, disks that are partitioned with MBR can have a size of up to 2 TiB.




Figure 7.1: MBR partitioning of the /dev/vdb storage device

The 2 TiB disk and partition size limit is now a common and restrictive limitation. Consequently, the legacy MBR scheme is superseded by the *GUID Partition Table (GPT)* partitioning scheme.

GPT Partition Scheme

For systems that run *Unified Extensible Firmware Interface (UEFI)* firmware, GPT is the standard for disk partitioning, and addresses the limitations of the MBR scheme. A GPT provides a maximum of 128 partitions. The GPT scheme allocates 64 bits for logical block addresses, to support partitions and disks of up to eight zebibytes (ZiB) or eight billion tebibytes (TiB).



Figure 7.2: GPT partitioning of the /dev/vdb storage device

GPT partitioning offers additional features and benefits over MBR. A GPT uses a *globally unique identifier* (GUID) to identify each disk and partition. A GPT makes the partition table redundant,

with the primary GPT at the head of the disk, and a backup secondary GPT at the end of the disk. A GPT uses a checksum to detect errors in the GPT header and partition table.

Manage Partitions

An administrator can use a *partition editor* program to change a disk's partitions, such as creating and deleting partitions, and changing partition types.

The standard partition editor on the command line in Red Hat Enterprise Linux is **parted**. You can use the **parted** partition editor with storage that uses either the MBR partitioning scheme or the GPT partitioning scheme.

The **parted** command takes as its first argument the device name that represents the entire storage device or disk to modify, followed by subcommands. The following example uses the **print** subcommand to display the partition table on the disk that is the **/dev/vda** block device (the first "virtualized I/O" disk detected by the system).

```
[root@host ~]# parted /dev/vda print
Model: Virtio Block Device (virtblk)
Disk /dev/vda: 53.7GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

Number  Start   End     Size    Type     File system  Flags
  1      1049kB  10.7GB  10.7GB  primary  xfs          boot
  2      10.7GB  53.7GB  42.9GB  primary  xfs
```

Use the **parted** command without a subcommand to open an interactive partitioning session.

```
[root@host ~]# parted /dev/vda
GNU Parted 3.4
Using /dev/vda
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) print
Model: Virtio Block Device (virtblk)
Disk /dev/vda: 53.7GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

Number  Start   End     Size    Type     File system  Flags
  1      1049kB  10.7GB  10.7GB  primary  xfs          boot
  2      10.7GB  53.7GB  42.9GB  primary  xfs

(parted) quit
[root@host ~]#
```

By default, the **parted** command displays sizes in powers of 10 (KB, MB, GB). You can change the unit size with the **unit** parameter, which accepts the following values:

- **s** for sector
- **B** for byte
- **MiB**, **GiB**, or **TiB** (powers of 2)

- **MB**, **GB**, or **TB** (powers of 10)

```
[root@host ~]# parted /dev/vda unit s print
Model: Virtio Block Device (virtblk)
Disk /dev/vda: 104857600s
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

Number  Start      End          Size         Type         File system  Flags
  1      2048s      20971486s    20969439s    primary      xfs          boot
  2      20971520s  104857535s   83886016s    primary      xfs
```

As shown in the previous example, you can also specify multiple subcommands (here, `unit` and `print`) on the same line.

Write the Partition Table on a New Disk

To partition a new drive, first write a disk label. The disk label indicates which partitioning scheme to use. Use `parted` to write an MBR disk label or a GPT disk label.

```
[root@host ~]# parted /dev/vdb mklabel msdos

[root@host ~]# parted /dev/vdb mklabel gpt
```



Warning

The `mklabel` subcommand wipes the existing partition table. Use the `mklabel` subcommand when the intent is to reuse the disk without regard to the existing data. If a new label moves the partition boundaries, then all data in existing file systems becomes inaccessible.

Create MBR Partitions

The following instructions create an MBR disk partition. Specify the disk device to create the partition on.

Run the `parted` command and specify the disk device name as an argument, to start in interactive mode. The session displays (`parted`) as a subcommand prompt.

```
[root@host ~]# parted /dev/vdb
GNU Parted 3.4
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted)
```

Use the `mkpart` subcommand to create a primary or extended partition.

```
(parted) mkpart
Partition type? primary/extended? primary
```

**Note**

If you need more than four partitions on an MBR-partitioned disk, then create three primary partitions and one extended partition. The extended partition serves as a container within which you can create multiple logical partitions.

Indicate the file-system type to create on the partition, such as `xf`s or `ext4`. This value is only a useful partition type label, and does not create the file system.

```
File system type? [ext2]? xfs
```

To list the supported file-system types, use the following command:

```
[root@host ~]# parted /dev/vdb help mkpart
...output omitted...
mkpart PART-TYPE [FS-TYPE] START END      make a partition

PART-TYPE is one of: primary, logical, extended
FS-TYPE is one of: udf, btrfs, nilfs2, ext4, ext3, ext2, f2fs, fat32, fat16,
hfsx, hfs+, hfs, jfs, swsusp, linux-swap(v1), linux-swap(v0), ntfs,
reiserfs, hp-ufs, sun-ufs, xfs, apfs2, apfs1, asfs, amufs5, amufs4, amufs3,
amufs2, amufs1, amufs0, amufs, affs7, affs6, affs5, affs4, affs3, affs2,
affs1, affs0, linux-swap, linux-swap(new), linux-swap(old)

'mkpart' makes a partition without creating a new file system on the
partition. FS-TYPE may be specified to set an appropriate partition
ID.
```

Specify the disk sector to start the new partition on.

```
Start? 2048s
```

The `s` suffix provides the value in sectors, or uses the `MiB`, `GiB`, `TiB`, `MB`, `GB`, or `TB` suffixes. The `parted` command defaults to the `MB` suffix. The `parted` command rounds provided values to satisfy disk constraints.

When the `parted` command starts, it retrieves the disk topology from the device, such as the disk physical block size. The `parted` command ensures that the start position that you provide correctly aligns the partition with the disk structure, to optimize performance. If the start position results in a misaligned partition, then the `parted` command displays a warning. With most disks, a start sector that is a multiple of 2048 is safe.

Specify the disk sector where the new partition should end, and exit `parted`. You can specify the end as a size or as an ending location.

```
End? 1000MB
(parted) quit
Information: You may need to update /etc/fstab.

[root@host ~]#
```

When you provide the end position, the `parted` command updates the partition table on the disk with the new partition details.

Run the `udevadm settle` command. This command waits for the system to detect the new partition and to create the associated device file in the `/dev` directory. The prompt returns when the task is done.

```
[root@host ~]# udevadm settle
```

As an alternative to interactive mode, you can create a partition in a single command:

```
[root@host ~]# parted /dev/vdb mkpart primary xfs 2048s 1000MB
```

Create GPT Partitions

The GPT scheme also uses the `parted` command to create partitions. Specify the disk device to create the partition on.

As the root user, execute the `parted` command and specify the disk device name as an argument.

```
[root@host ~]# parted /dev/vdb
GNU Parted 3.4
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted)
```

Use the `mkpart` subcommand to begin creating the partition. With the GPT scheme, each partition is given a name.

```
(parted) mkpart
Partition name? []? userdata
```

Indicate the file-system type to create on the partition, such as `xfs` or `ext4`. This value does not create the file system, but is a useful partition type label.

```
File system type? [ext2]? xfs
```

Specify the disk sector that the new partition starts on.

```
Start? 2048s
```

Specify the disk sector for the new partition to end, and exit `parted`. When you provide the end position, the `parted` command updates the GPT on the disk with the new partition details.

```
End? 1000MB
(parted) quit
Information: You may need to update /etc/fstab.

[root@host ~]#
```

Run the `udevadm settle` command. This command waits for the system to detect the new partition and to create the associated device file in the `/dev` directory. The prompt returns when the task is done.

```
[root@host ~]# udevadm settle
```

As an alternative to interactive mode, you can create a partition in a single command:

```
[root@host ~]# parted /dev/vdb mkpart userdata xfs 2048s 1000MB
```

Delete Partitions

The following instructions apply for both the MBR and GPT partitioning schemes. Specify the disk that contains the partition to remove.

Run the `parted` command with the disk device as the only argument.

```
[root@host ~]# parted /dev/vdb
GNU Parted 3.4
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted)
```

Identify the partition number of the partition to delete.

```
(parted) print
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 5369MB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name      Flags
  1      1049kB  1000MB  999MB   xfs          userdata
```

Delete the partition, and exit `parted`. The `rm` subcommand immediately deletes the partition from the partition table on the disk.

```
(parted) rm 1
(parted) quit
Information: You may need to update /etc/fstab.

[root@host ~]#
```

As an alternative to interactive mode, you can delete a partition in a single command:

```
[root@host ~]# parted /dev/vdb rm 1
```


Create File Systems

After a block device is created, the next step is to add a file system to it. Red Hat Enterprise Linux supports multiple file-system types, and XFS is the recommended default.

As the root user, use the `mkfs.xfs` command to apply an XFS file system to a block device. For an ext4 file system, use the `mkfs.ext4` command.

```
[root@host ~]# mkfs.xfs /dev/vdb1
meta-data=/dev/vdb1            isize=512    agcount=4, agsize=60992 blks
      =                       sectsz=512    attr=2, projid32bit=1
      =                       crc=1        finobt=1, sparse=1, rmapbt=0
      =                       reflink=1    bigtime=1 inobtcount=1
data      =                    bsize=4096   blocks=243968, imaxpct=25
      =                    sunit=0        swidth=0 blks
naming    =version 2           bsize=4096   ascii-ci=0, ftype=1
log       =internal log       bsize=4096   blocks=1566, version=2
      =                    sectsz=512    sunit=0 blks, lazy-count=1
realtime  =none               extsz=4096   blocks=0, rtextents=0
```

Mount File Systems

After you add the file system, the last step is to mount the file system to a directory in the directory structure. When you mount a file system to the directory hierarchy, user-space utilities can access or write files on the device.

Manually Mount File Systems

Use the `mount` command to manually attach a device to a *mount point* directory location. The `mount` command requires a device and a mount point, and can include file-system mount options. File-system options customize the behavior of the file system.

```
[root@host ~]# mount /dev/vdb1 /mnt
```

You also use the `mount` command to view currently mounted file systems, the mount points, and their options.

```
[root@host ~]# mount | grep vdb1
/dev/vdb1 on /mnt type xfs (rw,relatime,seclabel,attr2,inode64,noquota)
```

Persistently Mount File Systems

Manually mounting a file system is a good way to verify that a formatted device is accessible and is working as expected. However, when the server reboots, the system does not automatically mount the file system again.

To configure the system to automatically mount the file system during system boot, add an entry to the `/etc/fstab` file. This configuration file lists the file systems to mount at system boot.

The `/etc/fstab` file is a white-space-delimited file with six fields per line.

```
[root@host ~]# cat /etc/fstab

#
# /etc/fstab
# Created by anaconda on Thu Apr 5 12:05:19 2022
#
# 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.
#
# After editing this file, run 'systemctl daemon-reload' to update systemd
# units generated from this file.
#
UUID=a8063676-44dd-409a-b584-68be2c9f5570 / xfs defaults 0 0
UUID=7a20315d-ed8b-4e75-a5b6-24ff9e1f9838 /dbdata xfs defaults 0 0
```

The first field specifies the device. This example uses a UUID to specify the device. File systems create and store the UUID in the partition super block at creation time. Alternatively, you could use the device file, such as `/dev/vdb1`.

The second field is the directory mount point, from which the block device is accessible in the directory structure. The mount point must exist; if not, create it with the `mkdir` command.

The third field contains the file-system type, such as `xfs` or `ext4`.

The fourth field is the comma-separated list of options to apply to the device. `defaults` is a set of commonly used options. The `mount(8)` man page documents the other available options.

The fifth field is used by the `dump` command to back up the device. Other backup applications do not usually use this field.

The last field, the `fsck` order field, determines whether to run the `fsck` command at system boot to verify that the file systems are clean. The value in this field indicates the order in which `fsck` should run. For XFS file systems, set this field to `0`, because XFS does not use `fsck` to verify its file-system status. For `ext4` file systems, set it to `1` for the root file system, and to `2` for the other `ext4` file systems. By using this notation, the `fsck` utility processes the root file system first, and then verifies file systems on separate disks concurrently, and file systems on the same disk in sequence.



Note

An incorrect entry in `/etc/fstab` might render the machine non-bootable. Verify that an entry is valid by manually unmounting the new file system and then by using `mount /mountpoint` to read the `/etc/fstab` file, and remount the file system with that entry's mount options. If the `mount` command returns an error, then correct it before rebooting the machine.

Alternatively, use the `findmnt --verify` command to parse the `/etc/fstab` file for partition usability.

When you add or remove an entry in the `/etc/fstab` file, run the `systemctl daemon-reload` command, or reboot the server, to ensure that the `systemd` daemon loads and uses the new configuration.

```
[root@host ~]# systemctl daemon-reload
```

Red Hat recommends the use of UUIDs to persistently mount file systems, because block device names can change in certain scenarios, such as if a cloud provider changes the underlying storage layer of a virtual machine, or if disks are detected in a different order on a system boot. The block device file name might change, but the UUID remains constant in the file-system's super block.

Use the `lsblk --fs` command to scan the block devices that are connected to a machine and retrieve the file-system UUIDs.

```
[root@host ~]# lsblk --fs
NAME      FSTYPE  FSVER  LABEL      UUID                                FSAVAIL  FSUSE%  MOUNTPOINTS
vda
├─vda1
├─vda2 xfs          boot    49dd...75fdf    312M    37%    /boot
└─vda3 xfs          root    8a90...ce0da    4.8G    48%    /
```



References

`info parted` (*GNU Parted User Manual*)

`parted(8)`, `mkfs(8)`, `mount(8)`, `lsblk(8)`, and `fstab(5)` man pages

For more information, refer to the *Configuring and Managing File Systems* guide at https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/managing_file_systems/index

► Guided Exercise

Add Partitions, File Systems, and Persistent Mounts

Create a partition on a new storage device, format it with an XFS file system, configure it to mount at boot, and mount it for use.

Outcomes

- Use the `parted`, `mkfs.xfs`, and other commands to create a partition on a new disk, format it, and persistently mount it.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start storage-partitions
```

Instructions

- 1. Log in to `servera` as the `student` user and switch to the `root` user.

```
student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

- 2. Create an `msdos` disk label on the `/dev/vdb` device.

```
[root@servera ~]# parted /dev/vdb mklabel msdos
Information: You may need to update /etc/fstab.
```

- 3. Add a 1 GB primary partition. For correct alignment, start the partition at the 2048 sector. Set the partition file-system type to XFS.

- 3.1. Use `parted` interactive mode to create the partition.

```
[root@servera ~]# parted /dev/vdb
GNU Parted 3.4
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) mkpart
```

```
Partition type? primary/extended? primary
File system type? [ext2]? xfs
Start? 2048s
End? 1001MB
(parted) quit
Information: You may need to update /etc/fstab.
```

Because the partition starts at the 2048 sector, the previous command sets the end position to 1001 MB to get a partition size of 1000 MB (1 GB).

Alternatively, you can perform the same operation with the following non-interactive command: `parted /dev/vdb mkpart primary xfs 2048s 1001 MB`

- 3.2. Verify your work by listing the partitions on the `/dev/vdb` device.

```
[root@servera ~]# parted /dev/vdb print
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 5369MB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

Number   Start    End      Size    Type    File system  Flags
  1       1049KB   1001MB   1000MB   primary
```

- 3.3. Run the `udevadm settle` command. This command waits for the system to register the new partition, and returns when it is done.

```
[root@servera ~]# udevadm settle
```

- ▶ 4. Format the new partition with the XFS file system.

```
[root@servera ~]# mkfs.xfs /dev/vdb1
meta-data=/dev/vdb1            isize=512    agcount=4, agsize=61056 blks
                               =               sectsz=512   attr=2, projid32bit=1
                               =               crc=1     finobt=1, sparse=1, rmapbt=0
                               =               reflink=1  bigtime=1 inobtcount=1
data      =                     bsize=4096   blocks=244224, imaxpct=25
                               =               sunit=0    swidth=0 blks
naming    =version 2           bsize=4096   ascii-ci=0, ftype=1
log       =internal log        bsize=4096   blocks=1566, version=2
                               =               sectsz=512   sunit=0 blks, lazy-count=1
realtime  =none                extsz=4096   blocks=0, rtextents=0
```

- ▶ 5. Configure the new file system to mount to the `/archive` directory persistently.

- 5.1. Create the `/archive` directory.

```
[root@servera ~]# mkdir /archive
```

- 5.2. Discover the UUID of the `/dev/vdb1` device. The UUID in the output is probably different on your system.

```
[root@servera ~]# lsblk --fs /dev/vdb
NAME      FSTYPE FSVER LABEL UUID                                 FSAVAIL FSUSE%
MOUNTPOINTS
vdb
└─vdb1 xfs                                     881e856c-37b1-41e3-b009-ad526e46d987
```

- 5.3. Add an entry to the `/etc/fstab` file. Replace the UUID with the one that you discovered from the previous step.

```
...output omitted...
UUID=881e856c-37b1-41e3-b009-ad526e46d987 /archive xfs defaults 0 0
```

- 5.4. Update the `systemd` daemon for the system to register the new `/etc/fstab` file configuration.

```
[root@servera ~]# systemctl daemon-reload
```

- 5.5. Mount the new file system with the new entry in the `/etc/fstab` file.

```
[root@servera ~]# mount /archive
```

- 5.6. Verify that the new file system is mounted on the `/archive` directory.

```
[root@servera ~]# mount | grep /archive
/dev/vdb1 on /archive type xfs
(rw,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,noquota)
```

- ▶ 6. Reboot `servera`. After the server rebooted, log in and verify that the `/dev/vdb1` device is mounted on the `/archive` directory. When done, log out from `servera`.

- 6.1. Reboot `servera`.

```
[root@servera ~]# systemctl reboot
Connection to servera closed by remote host.
Connection to servera closed.
[student@workstation ~]$
```

- 6.2. Wait for `servera` to reboot and log in as the `student` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
```

- 6.3. Verify that the `/dev/vdb1` device is mounted on the `/archive` directory.

```
[student@servera ~]$ mount | grep /archive
/dev/vdb1 on /archive type xfs
(rw,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,noquota)
```

6.4. Return to the `workstation` machine as the `student` user.

```
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish storage-partitions
```

Manage Swap Space

Objectives

- Create and manage swap spaces to supplement physical memory.

Swap Space Concepts

A *swap space* is an area of a disk under the control of the Linux kernel memory management subsystem. The kernel uses swap space to supplement the system RAM by holding inactive pages in memory. A system's *virtual memory* encompasses the combined system RAM and swap space.

When the memory usage on a system exceeds a defined limit, the kernel searches through RAM to look for idle memory pages that are assigned to processes. The kernel writes the idle pages to the swap area and reassigns the RAM pages to other processes. If a program requires access to a page on disk, then the kernel locates another idle page of memory, writes it to disk, and recalls the needed page from the swap area.

Because swap areas are on disk, swap is slow when compared with RAM. Although swap space augments system RAM, do not consider swap space as a sustainable solution for insufficient RAM for your workload.

Swap Space Calculation

Administrators should size the swap space based on the memory workload on the system. Application vendors sometimes provide recommendations for calculating swap space. The following table provides guidance based on the total physical memory.

RAM and Swap Space Recommendations

RAM	Swap space	Swap space if allowing for hibernation
2 GB or less	Twice the RAM	Three times the RAM
Between 2 GB and 8 GB	Same as RAM	Twice the RAM
Between 8 GB and 64 GB	At least 4 GB	1.5 times the RAM
More than 64 GB	At least 4 GB	Hibernation is not recommended

The laptop and desktop hibernation function uses the swap space to save the RAM contents before powering off the system. When you turn the system back on, the kernel restores the RAM contents from the swap space and does not need a complete boot. For those systems, the swap space must be greater than the amount of RAM.

The Knowledgebase article in References at the end of this section gives more guidance about sizing the swap space.

Create Swap Space

To create a swap space, perform the following steps:

- Create a partition with a file-system type of `linux-swap`.
- Place a swap signature on the device.

Create a Swap Partition

Use the `parted` command to create a partition of the appropriate size and set its file-system type to `linux-swap`. In the past, tools determined from the partition file-system type whether to activate the device; however, that requirement is no longer the case. Even though utilities no longer use the partition file-system type, administrators can determine the partition's purpose from that type.

The following example creates a 256 MB partition.

```
[root@host ~]# parted /dev/vdb
GNU Parted 3.4
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) print
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 5369MB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name  Flags
  1      1049kB  1001MB  1000MB                data

(parted) mkpart
Partition name? []? swap1
File system type? [ext2]? linux-swap
Start? 1001MB
End? 1257MB
(parted) print
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 5369MB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name  Flags
  1      1049kB  1001MB  1000MB                data
  2      1001MB  1257MB  256MB    linux-swap(v1)  swap1

(parted) quit
Information: You may need to update /etc/fstab.

[root@host ~]#
```

After creating the partition, run the `udevadm settle` command. This command waits for the system to detect the new partition and to create the associated device file in the `/dev` directory. The command returns only when it is finished.

```
[root@host ~]# udevadm settle
```

Format Swap Space

The `mkswap` command applies a swap signature to the device. Unlike other formatting utilities, the `mkswap` command writes a single block of data at the beginning of the device, and leaves the rest of the device unformatted so that the kernel can use it for storing memory pages.

```
[root@host ~]# mkswap /dev/vdb2
Setting up swspace version 1, size = 244 MiB (255848448 bytes)
no label, UUID=39e2667a-9458-42fe-9665-c5c854605881
```

Activate Swap Space

You can use the `swapon` command to activate a formatted swap space.

Use `swapon` with the device as a parameter, or use `swapon -a` to activate all the listed swap spaces in the `/etc/fstab` file. Use the `swapon --show` and `free` commands to inspect the available swap spaces.

```
[root@host ~]# free
              total        used        free      shared  buff/cache   available
Mem:          1873036       134688       1536436         16748         201912       1576044
Swap:           0              0              0

[root@host ~]# swapon /dev/vdb2
[root@host ~]# free
              total        used        free      shared  buff/cache   available
Mem:          1873036       135044       1536040         16748         201952       1575680
Swap:       249852              0       249852
```

You can deactivate a swap space with the `swapoff` command. If pages are written to the swap space, then the `swapoff` command tries to move those pages to other active swap spaces or back into memory. If the `swapoff` command cannot write data to other places, then it fails with an error, and the swap space stays active.

Activate Swap Space Persistently

Create an entry in the `/etc/fstab` file to ensure an active swap space at system boot. The following example shows a typical line in the `/etc/fstab` file based on the previously created swap space.

```
UUID=39e2667a-9458-42fe-9665-c5c854605881    swap    swap    defaults    0 0
```

The example uses the UUID as the first field. When you format the device, the `mkswap` command displays that UUID. If you lost the output of `mkswap`, then use the `lsblk --fs` command. As an alternative, you can use the device name in the first field.

The second field is typically reserved for the mount point. However, for swap devices, which are not accessible through the directory structure, this field takes the `swap` placeholder value. The `fstab(5)` man page uses a `none` placeholder value; however, a `swap` value gives more informative error messages if something goes wrong.

The third field is the file-system type. The file-system type for swap space is `swap`.

The fourth field is for options. The example uses the `defaults` option. The `defaults` option includes the `auto` mount option, which activates the swap space automatically at system boot.

The final two fields are the `dump` flag and the `fsck` order. Swap spaces do not require backing up or file-system checking, and so these fields should be set to zero.

When you add or remove an entry in the `/etc/fstab` file, run the `systemctl daemon-reload` command, or reboot the server, for `systemd` to register the new configuration.

```
[root@host ~]# systemctl daemon-reload
```

Set Swap Space Priority

By default, the system uses swap spaces in priority order. The kernel uses the highest priority swap space until it is full, and then starts to use the low priority swap space.

The default priority for swap spaces is low, and new swap spaces are lower priority than older swap spaces. When swap spaces have the same priority, the kernel writes to them in a round-robin fashion.

To set the priority, use the `pri` option in the `/etc/fstab` file. The kernel uses the swap space with the highest priority first. The default priority is `-2`.

The following example shows three defined swap spaces in the `/etc/fstab` file. The kernel uses the last entry first, because its priority is set to 10. When that space is full, it uses the second entry, because its priority is set to 4. Finally, it uses the first entry, which has a default priority of `-2`.

```
UUID=af30cbb0-3866-466a-825a-58889a49ef33  swap  swap  defaults  0 0
UUID=39e2667a-9458-42fe-9665-c5c854605881  swap  swap  pri=4     0 0
UUID=fbd7fa60-b781-44a8-961b-37ac3ef572bf   swap  swap  pri=10    0 0
```

Use the `swapon --show` command to display the swap space priorities.



References

`mkswap(8)`, `swapon(8)`, `swapoff(8)`, `mount(8)`, and `parted(8)` man pages

Knowledgebase: What Is the Recommended Swap Size for Red Hat Platforms?

<https://access.redhat.com/solutions/15244>

► Guided Exercise

Manage Swap Space

Create and format a partition as a swap space, and activate it persistently.

Outcomes

- Create a partition and a swap space on a disk by using the GPT partitioning scheme.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start storage-swap
```

Instructions

- 1. Log in to `servera` as the `student` user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

- 2. Inspect the `/dev/vdb` disk. The disk already has a partition table and uses the GPT partitioning scheme. Also, it has an existing 1 GB partition.

```
[root@servera ~]# parted /dev/vdb print
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 5369MB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name  Flags
  1      1049kB  1001MB  1000MB                data
```

- 3. Add a new partition of 500 MB for use as a swap space. Set the partition type to `linux-swap`.
- 3.1. Create the `myswap` partition. Because the disk uses the GPT partitioning scheme, you must give a name to the partition. Notice that the start position, 1001 MB, is the end of the existing first partition. The `parted` command ensures that the new partition immediately follows the previous one, without any gap. Because the

partition starts at the 1001 MB position, the command sets the end position to 1501 MB to get a partition size of 500 MB.

```
[root@servera ~]# parted /dev/vdb mkpart myswap linux-swap \
1001MB 1501MB
Information: You may need to update /etc/fstab.
```

- 3.2. Verify your work by listing the partitions on the /dev/vdb disk. The size of the new partition is not exactly 500 MB. The difference in size is because the parted command must align the partition with the disk layout.

```
[root@servera ~]# parted /dev/vdb print
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 5369MB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name  Flags
  1      1049kB  1001MB  1000MB                data
  2      1001MB  1501MB  499MB                myswap swap
```

- 3.3. Run the udevadm settle command. This command waits for the system to register the new partition, and returns when it is done.

```
[root@servera ~]# udevadm settle
```

- ▶ 4. Initialize the new partition as a swap space.

```
[root@servera ~]# mkswap /dev/vdb2
Setting up swspace version 1, size = 476 MiB (499118080 bytes)
no label, UUID=cb7f71ca-ee82-430e-ad4b-7dda12632328
```

- ▶ 5. Enable the new swap space.

- 5.1. Verify that creating and initializing the swap space does not yet enable it for use.

```
[root@servera ~]# swapon --show
```

- 5.2. Enable the new swap space.

```
[root@servera ~]# swapon /dev/vdb2
```

- 5.3. Verify that the new swap space is now available.

```
[root@servera ~]# swapon --show
NAME      TYPE      SIZE USED PRI0
/dev/vdb2 partition 476M   0B   -2
```

- 5.4. Disable the swap space.

```
[root@servera ~]# swapoff /dev/vdb2
```

5.5. Confirm that the swap space is disabled.

```
[root@servera ~]# swapon --show
```

► 6. Enable the new swap space at system boot.

6.1. Use the `lsblk` command with the `--fs` option to discover the UUID of the `/dev/vdb2` device. The UUID in the output is different on your system.

```
[root@servera ~]# lsblk --fs /dev/vdb2
```

NAME	FSTYPE	FSVER	LABEL	UUID	FSAVAIL	FSUSE%
MOUNTPOINTS						
vdb2	swap	1		762735cb-a52a-4345-9ed0-e3a68aa8bb97		

6.2. Add an entry to the `/etc/fstab` file. In the following command, replace the UUID with the one that you discovered from the previous step.

```
...output omitted...
UUID=762735cb-a52a-4345-9ed0-e3a68aa8bb97 swap swap defaults 0 0
```

6.3. Update the `systemd` daemon for the system to register the new `/etc/fstab` file configuration.

```
[root@servera ~]# systemctl daemon-reload
```

6.4. Enable the swap space by using the entry in the `/etc/fstab` file.

```
[root@servera ~]# swapon -a
```

6.5. Verify that the new swap space is enabled.

```
[root@servera ~]# swapon --show
```

NAME	TYPE	SIZE	USED	PRIO
/dev/vdb2	partition	476M	0B	-2

► 7. Reboot the `servera` machine. After the server reboots, log in and verify that the swap space is enabled. When done, log out from `servera`.

7.1. Reboot the `servera` machine.

```
[root@servera ~]# systemctl reboot
Connection to servera closed by remote host.
Connection to servera closed.
[student@workstation ~]$
```

7.2. Wait for `servera` to reboot and log in as the `student` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
```

7.3. Verify that the swap space is enabled.

```
[student@servera ~]$ swapon --show
NAME      TYPE      SIZE USED PRI0
/dev/vdb2 partition 476M   0B   -2
```

7.4. Return to the `workstation` machine as the `student` user.

```
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish storage-swap
```

► Lab

Manage Basic Storage

Create several partitions on a new disk, and format some with file systems and mount them, and activate others as swap spaces.

Outcomes

- Display and create partitions with the `parted` command.
- Create file systems on partitions and persistently mount them.
- Create swap spaces and activate them at boot.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start storage-review
```

Instructions

1. The `serverb` machine has several unused disks. On the first unused disk, create a GPT partition label and a 2 GB GPT partition named `backup`.
Because it is difficult to set an exact size, a size between 1.8 GB and 2.2 GB is acceptable.
Configure the `backup` partition to host an XFS file system.
2. Format the 2 GB `backup` partition with an XFS file system and persistently mount it to the `/backup` directory by using the UUID.
3. On the same disk, create two 512 MB GPT partitions with the `swap1` and `swap2` names.
A size between 460 MB and 564 MB is acceptable.
Configure the file-system types of the partitions to host swap spaces.
4. Initialize the two 512 MB partitions as swap spaces, and configure them to activate at boot. Set the swap space on the `swap2` partition to be preferred over the other. Note that 512 MB is approximately equivalent to 488 MiB.
5. To verify your work, reboot the `serverb` machine. Confirm that the system automatically mounts the first partition to the `/backup` directory. Also, confirm that the system activates the two swap spaces.

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.


```
[student@workstation ~]$ lab grade storage-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish storage-review
```

► Solution

Manage Basic Storage

Create several partitions on a new disk, and format some with file systems and mount them, and activate others as swap spaces.

Outcomes

- Display and create partitions with the `parted` command.
- Create file systems on partitions and persistently mount them.
- Create swap spaces and activate them at boot.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start storage-review
```

Instructions

1. The `serverb` machine has several unused disks. On the first unused disk, create a GPT partition label and a 2 GB GPT partition named `backup`.
Because it is difficult to set an exact size, a size between 1.8 GB and 2.2 GB is acceptable.
Configure the `backup` partition to host an XFS file system.
 - 1.1. Log in to `serverb` as the `student` user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
[student@serverb ~]$ sudo -i
[sudo] password for student: student
[root@serverb ~]#
```

- 1.2. Identify the unused disks. The first unused disk, `/dev/vdb`, does not have any partitions.

```
[root@serverb ~]# lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINTS
vda         252:0    0   10G  0 disk
└─vda1      252:1    0    1M  0 part
└─vda2      252:2    0  200M  0 part /boot/efi
└─vda3      252:3    0  500M  0 part /boot
└─vda4      252:4    0   9.3G  0 part /
```

vdb	252:16	0	5G	0	disk
vdc	252:32	0	5G	0	disk
vdd	252:48	0	5G	0	disk

1.3. Confirm that the `/dev/vdb` disk has no label.

```
[root@serverb ~]# parted /dev/vdb print
Error: /dev/vdb: unrecognised disk label
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 5369MB
Sector size (logical/physical): 512B/512B
Partition Table: unknown
Disk Flags:
```

1.4. Define the GPT partitioning scheme.

```
[root@serverb ~]# parted /dev/vdb mklabel gpt
Information: You may need to update /etc/fstab.
```

1.5. Create the 2 GB backup partition with an xfs file-system type. Start the partition at sector 2048.

```
[root@serverb ~]# parted /dev/vdb mkpart backup xfs 2048s 2GB
Information: You may need to update /etc/fstab.
```

1.6. Confirm the creation of the backup partition.

```
[root@serverb ~]# parted /dev/vdb print
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 5369MB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
```

Number	Start	End	Size	File system	Name	Flags
1	1049KB	2000MB	1999MB		backup	

1.7. Run the `udevadm settle` command. This command waits for the system to detect the new partition and to create the `/dev/vdb1` device file.

```
[root@serverb ~]# udevadm settle
```

2. Format the 2 GB backup partition with an XFS file system and persistently mount it to the `/backup` directory by using the UUID.

2.1. Format the `/dev/vdb1` partition.

```
[root@serverb ~]# mkfs.xfs /dev/vdb1
meta-data=/dev/vdb1             isize=512    agcount=4, agsize=121984 blks
=                               sectsz=512   attr=2, projid32bit=1
=                               crc=1        finobt=1, sparse=1, rmapbt=0
```

```
=
data      =                                reflink=1    bigtime=1 inobtcount=1
          =                                bsize=4096    blocks=487936, imaxpct=25
          =                                sunit=0       swidth=0 blks
naming    =version 2                      bsize=4096    ascii-ci=0, ftype=1
log        =internal log                  bsize=4096    blocks=2560, version=2
          =                                sectsz=512    sunit=0 blks, lazy-count=1
realtime   =none                          extsz=4096    blocks=0, rtextents=0
```

2.2. Create the `/backup` mount point.

```
[root@serverb ~]# mkdir /backup
```

2.3. Before adding the new file system to the `/etc/fstab` file, retrieve its UUID. The UUID on your system might be different.

```
[root@serverb ~]# lsblk --fs /dev/vdb1
NAME FSTYPE FSVER LABEL UUID                                FSAVAIL FSUSE%
MOUNTPOINTS
vdb1 xfs                                           f74ed805-b1fc-401a-a5ee-140f97c6757d
```

2.4. Edit the `/etc/fstab` file and define the new file system.

```
[root@serverb ~]# vim /etc/fstab
...output omitted...
UUID=f74ed805-b1fc-401a-a5ee-140f97c6757d /backup xfs defaults 0 0
```

2.5. Force the `systemd` daemon to reread the `/etc/fstab` file.

```
[root@serverb ~]# systemctl daemon-reload
```

2.6. Manually mount the `/backup` directory to verify your work. Confirm that the mount is successful.

```
[root@serverb ~]# mount /backup
[root@serverb ~]# mount | grep /backup
/dev/vdb1 on /backup type xfs
(rw,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,noquota)
```

3. On the same disk, create two 512 MB GPT partitions with the `swap1` and `swap2` names. A size between 460 MB and 564 MB is acceptable. Configure the file-system types of the partitions to host swap spaces.

3.1. Retrieve the end position of the first partition by displaying the current partition table on the `/dev/vdb` disk. In the next step, you use that value as the start of the `swap1` partition.

```
[root@serverb ~]# parted /dev/vdb print
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 5369MB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
```

Disk Flags:

Number	Start	End	Size	File system	Name	Flags
1	1049kB	2000MB	1999MB	xfs	backup	

- 3.2. Create the first 512 MB GPT partition named `swap1`. Set its type to `linux-swap`. Use the end position of the first partition as the starting point. The end position is 2000 MB + 512 MB = 2512 MB.

```
[root@serverb ~]# parted /dev/vdb mkpart swap1 linux-swap 2000M 2512M
Information: You may need to update /etc/fstab.
```

- 3.3. Create the second 512 MB GPT partition named `swap2`. Set its type to `linux-swap`. Use the end position of the previous partition as the starting point: 2512M. The end position is 2512 MB + 512 MB = 3024 MB.

```
[root@serverb ~]# parted /dev/vdb mkpart swap2 linux-swap 2512M 3024M
Information: You may need to update /etc/fstab.
```

- 3.4. Display the partition table to verify your work.

```
[root@serverb ~]# parted /dev/vdb print
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 5369MB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
```

Number	Start	End	Size	File system	Name	Flags
1	1049kB	2000MB	1999MB	xfs	backup	
2	2000MB	2512MB	513MB		swap1	swap
3	2512MB	3024MB	512MB		swap2	swap

- 3.5. Run the `udevadm settle` command. The command waits for the system to register the new partitions and to create the device files.

```
[root@serverb ~]# udevadm settle
```

4. Initialize the two 512 MB partitions as swap spaces, and configure them to activate at boot. Set the swap space on the `swap2` partition to be preferred over the other. Note that 512 MB is approximately equivalent to 488 MiB.
 - 4.1. Use the `mkswap` command to initialize the swap partitions. Note the UUIDs of the two swap spaces, because you use that information in the next step. If you clear the `mkswap` output, then use the `lsblk --fs` command to retrieve the UUIDs.

```
[root@serverb ~]# mkswap /dev/vdb2
Setting up swapspace version 1, size = 489 MiB (512749568 bytes)
no label, UUID=87976166-4697-47b7-86d1-73a02f0fc803
[root@serverb ~]# mkswap /dev/vdb3
Setting up swapspace version 1, size = 488 MiB (511700992 bytes)
no label, UUID=4d9b847b-98e0-4d4e-9ef7-dfaaf736b942
```

- 4.2. Edit the `/etc/fstab` file and define the new swap spaces. To set the swap space on the `swap2` partition to be preferred over the `swap1` partition, give the `swap2` partition a higher priority with the `pri` option.

```
[root@serverb ~]# vim /etc/fstab
...output omitted...
UUID=a3665c6b-4bfb-49b6-a528-74e268b058dd /backup xfs defaults 0 0
UUID=87976166-4697-47b7-86d1-73a02f0fc803 swap swap pri=10 0 0
UUID=4d9b847b-98e0-4d4e-9ef7-dfaaf736b942 swap swap pri=20 0 0
```

- 4.3. Force the `systemd` daemon to reread the `/etc/fstab` file.

```
[root@serverb ~]# systemctl daemon-reload
```

- 4.4. Activate the new swap spaces. Verify the correct activation of the swap spaces.

```
[root@serverb ~]# swapon -a
[root@serverb ~]# swapon --show
NAME      TYPE      SIZE USED PRI0
/dev/vdb2 partition 489M  0B   10
/dev/vdb3 partition 488M  0B   20
```

5. To verify your work, reboot the `serverb` machine. Confirm that the system automatically mounts the first partition to the `/backup` directory. Also, confirm that the system activates the two swap spaces.

- 5.1. Reboot `serverb`.

```
[root@serverb ~]# systemctl reboot
Connection to serverb closed by remote host.
Connection to serverb closed.
[student@workstation ~]$
```

- 5.2. Wait for `serverb` to boot, and then log in as the `student` user.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
[student@serverb ~]$
```

- 5.3. Verify that the system automatically mounts the `/dev/vdb1` partition to the `/backup` directory.

```
[student@serverb ~]$ mount | grep /backup
/dev/vdb1 on /backup type xfs (rw,relatime,seclabel,attr2,inode64,noquota)
```

- 5.4. Verify that the system activates both swap spaces.

```
[student@serverb ~]$ swapon --show
NAME      TYPE      SIZE USED PRI0
/dev/vdb2 partition 489M  0B   10
/dev/vdb3 partition 488M  0B   20
```

5.5. Return to the `workstation` machine as the `student` user.

```
[student@serverb ~]$ exit
logout
Connection to serverb closed.
[student@workstation ~]$
```

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade storage-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish storage-review
```

Summary

- The `parted` command adds, modifies, and removes partitions on disks with the MBR or the GPT partitioning scheme.
- The `mkfs.xfs` command creates XFS file systems on disk partitions.
- The `/etc/fstab` file contains devices that must be persistently mounted.
- The `mkswap` command initializes swap spaces.

Chapter 8

Manage Storage Stack

Goal

Create and manage logical volumes that contain file systems or swap spaces from the command line.

Objectives

- Describe logical volume manager components and concepts, and implement LVM storage and display LVM component information.
- Analyze the multiple storage components that make up the layers of the storage stack.

Sections

- Create and Extend Logical Volumes (and Guided Exercise)
- Manage Layered Storage (and Guided Exercise)

Lab

- Manage Storage Stack

Create and Extend Logical Volumes

Objectives

- Describe logical volume manager components and concepts, and implement LVM storage and display LVM component information.

Logical Volume Manager Overview

Use the *Logical Volume Manager (LVM)* system to create logical storage volumes as a layer on the physical storage. This storage system provides greater flexibility than using physical storage directly. LVM hides the hardware storage configuration from the software, and enables you to resize volumes without stopping applications or unmounting file systems. LVM provides comprehensive command-line tools to manage storage.

Physical devices

Logical volumes use physical devices for storing data. These devices might be disk partitions, whole disks, RAID arrays, or SAN disks. You must initialize the device as an LVM physical volume. An LVM physical volume must use the entire physical device.

Physical Volumes (PVs)

LVM uses the underlying physical device as the LVM physical volume. LVM tools segment the physical volumes into *Physical Extents (PEs)* to form small chunks of data that act as the smallest storage block on a PV.

Volume Groups (VGs)

Volume groups are storage pools that are made from one or more PVs. It is the functional equivalent of a whole disk in physical storage. A PV must be allocated only to a single VG. LVM sets the PE size automatically, although it is possible to specify it. A VG might consist of unused space and several logical volumes.

Logical Volumes (LVs)

Logical volumes are created from free physical extents in a VG, and are provided as the storage device for applications, users, and operating systems. LVs are a collection of *Logical Extents (LEs)*, which map to physical extents. By default, each LE gets mapped to one PE. Setting specific LV options changes this mapping; for example, mirroring causes each LE to map to two PEs.

Logical Volume Manager Workflow

Creating LVM storage requires building structures in a logical workflow.

- Determine the physical devices that are used for creating physical volumes, and initialize these devices as LVM physical volumes.
- Create a volume group from multiple physical volumes.
- Create the logical volumes from the available space in the volume group.
- Format the logical volume with a file system and mount it, or activate it as swap space, or pass the raw volume to a database or storage server for advanced structures.




Figure 8.1: Logical Volume Manager workflow

**Note**

The examples here use a /dev/vdb device name and its storage partitions. The device names on your classroom system might be different. Use the `lsblk`, `blkid`, or `cat /proc/partitions` commands to identify your system's devices.

Build LVM Storage

Creating a logical volume involves creating physical device partitions, physical volumes, and volume groups. After creating an LV, format the volume and mount it to access it as storage.

Prepare Physical Devices

Partitioning is optional when already present. Use the `parted` command to create a partition on the physical device. Set the physical device to the Linux LVM partition type. Use the `udevadm settle` command to register the new partition with the kernel.

```
[root@host ~]# parted /dev/vdb mklabel gpt mkpart primary 1MiB 769MiB
...output omitted...
[root@host ~]# parted /dev/vdb mkpart primary 770MiB 1026MiB
[root@host ~]# parted /dev/vdb set 1 lvm on
[root@host ~]# parted /dev/vdb set 2 lvm on
[root@host ~]# udevadm settle
```

Create Physical Volumes

Use the `pvcreeate` command to label the physical partition as an LVM physical volume. Label multiple devices simultaneously by using space-delimited device names as arguments to the `pvcreeate` command. This example labels the /dev/vdb1 and /dev/vdb2 devices as PVs that are ready for creating volume groups.

```
[root@host ~]# pvcreate /dev/vdb1 /dev/vdb2
Physical volume "/dev/vdb1" successfully created.
Physical volume "/dev/vdb2" successfully created.
Creating devices file /etc/lvm/devices/system.devices
```

Create a Volume Group

The `vgcreate` command builds one or more physical volumes into a volume group. The first argument is a volume group name, followed by one or more physical volumes to allocate to this VG. This example creates the `vg01` VG by using the `/dev/vdb1` and `/dev/vdb2` PVs.

```
[root@host ~]# vgcreate vg01 /dev/vdb1 /dev/vdb2
Volume group "vg01" successfully created
```

Create a Logical Volume

The `lvcreate` command creates a logical volume from the available PEs in a volume group. Use the `lvcreate` command to set the LV name and size, and the VG name to contain this logical volume. This example creates the `lv01` LV with 300 MiB in the `vg01` VG.

```
[root@host ~]# lvcreate -n lv01 -L 300M vg01
Logical volume "lv01" created.
```

This command might fail if the volume group does not have enough free physical extents. The LV size rounds up to the next PE size value when the size does not exactly match.

The `lvcreate` command `-L` option requires sizes in bytes, mebibytes (binary megabytes, 1048576 bytes), and gibibytes (binary gigabytes), or similar. The lowercase `-l` requires sizes that are specified as a number of physical extents. The following commands are two choices for creating the same LV with the same size:

- `lvcreate -n lv01 -L 128M vg01` : create an LV of size 128 MiB, rounded to the next PE.
- `lvcreate -n lv01 -l 32 vg01` : create an LV of size 32 PEs at 4 MiB each, total 128 MiB.

Create a Logical Volume with Deduplication and Compression

RHEL 9 uses an LVM *Virtual Data Optimizer (VDO)* implementation for managing VDO volumes. The previous Python-based VDO management tools are still available but are no longer needed.

The VDO provides inline block-level deduplication, compression, and thin provisioning for storage. Configure a VDO volume to use up to 256 TB of physical storage. Manage VDO as a type of LVM logical volume (LVs), similar to LVM thinly provisioned volumes. An LVM VDO is composed of two logical volumes:

VDO pool LV

This LV stores, deduplicates, compresses data, and sets the size of the VDO volume that is backed by the physical device. VDO is deduplicated and compresses each VDO LV separately, because each VDO pool LV can hold only one VDO LV.

VDO LV

A virtual device is provisioned on top of the VDO pool LV, and sets the logical size of the VDO volume to store the data before deduplication and compression occur.

LVM VDO presents the deduplicated storage as a regular logical volume (LV). The VDO volume can be formatted with standard file systems, or shared as a block device, or used to build other storage layers, the same as any normal logical volume.

To use VDO deduplication and compression, install the `vdo` and `kmod-kvdo` packages.

```
[root@host ~]# dnf install vdo kmod-kvdo
```

Verify that the selected LVM volume group has enough free storage capacity. Use the `lvcreate` command with the `--type vdo` parameter to create a VDO LV.

```
[root@host ~]# lvcreate --type vdo --name vdo-lv01 --size 5G vg01
Logical blocks defaulted to 523108 blocks.
The VDO volume can address 2 GB in 1 data slab.
It can grow to address at most 16 TB of physical storage in 8192 slabs.
If a larger maximum size might be needed, use bigger slabs.
Logical volume "vdo-lv01" created.
```

Create a File System on the Logical Volume

Specify the logical volume by using either the `/dev/vgname/lvname` traditional name, or the `/dev/mapper/vgname-lvname` kernel device mapper name.

Use the `mkfs` command to create a file system on the new logical volume.

```
[root@host ~]# mkfs -t xfs /dev/vg01/vdo-lv01
...output omitted...
```

Create a mount point by using the `mkdir` command.

```
[root@host ~]# mkdir /mnt/data
```

To make the file system available persistently, add an entry to the `/etc/fstab` file.

```
/dev/vg01/vdo-lv01 /mnt/data xfs defaults 0 0
```

Mount the LV by using the `mount` command.

```
[root@host ~]# mount /mnt/data/
```



Note

You can mount a logical volume by name or by UUID, because LVM parses the PVs by looking for the UUID. This behavior is successful even when the VG was created by using a name, because the PV always contains a UUID.

Display LVM Component Status

LVM provides various utilities to display the status information of PV, VG, and LV. Use the `pvdiskdisplay`, `vgdisplay`, and `lvdisplay` commands to show the status information of the LVM components.

The associated `pvs`, `vgs`, and `lvs` commands are commonly used and show a subset of the status information, with one line for each entity.

Display Physical Volume Information

The `pvdiskdisplay` command displays information about the PVs. Use the command without arguments to list information of all PVs on the system. Providing the name of the PV as the argument to the command shows the information that is specific to the PV.

```
[root@host ~]# pvdiskdisplay /dev/vdb1
--- Physical volume ---
PV Name                /dev/vdb1
VG Name                vg01
PV Size                731.98 MiB / not usable 3.98 MiB
Allocatable            yes
PE Size                4.00 MiB
Total PE               182
Free PE                107
Allocated PE           75
PV UUID                zP0gD9-NxTV-Qtoi-yfQD-TGpL-0Yj0-wExh2N
```

- ❶ PV Name shows the device name.
- ❷ VG Name shows the volume group where the PV is allocated.
- ❸ PV Size shows the physical size of the PV, including unusable space.
- ❹ PE Size shows the physical extent size.
- ❺ Free PE shows the PE size available in the VG to create new LVs or extend existing LVs.

Display Volume Group Information

The `vgdisplay` command shows the information about volume groups. To list information about all VGs, use the command without arguments. Provide the name of the VG as an argument to list information that is specific to the VG.

```
[root@host ~]# vgdisplay vg01
--- Volume group ---
VG Name                vg01
System ID
Format                lvm2
Metadata Areas         2
Metadata Sequence No   2
VG Access              read/write
VG Status              resizable
MAX LV                 0
Cur LV                1
Open LV                1
```

Max PV	0	
Cur PV	2	
Act PV	2	
VG Size	1012.00 MiB	2
PE Size	4.00 MiB	
Total PE	253	3
Alloc PE / Size	75 / 300.00 MiB	
Free PE / Size	178 / 712.00 MiB	4
VG UUID	jK5M1M-YvIk-kxU2-bxMS-dNjQ-Bs3L-DRLJNc	

- 1 VG Name shows the name of the volume group.
- 2 VG Size displays the total size of the storage pool that is available for LV allocation.
- 3 Total PE shows the total size of PE units.
- 4 Free PE / Size shows the available space in the VG to create or extend LVs.

Display Logical Volume Information

The `lvdisplay` command displays information about logical volumes. Use the command without arguments to list information of all LVs. Providing the name of the LV as an argument displays the information that is specific to the LV.

```
[root@host ~]# lvdisplay /dev/vg01/lv01
--- Logical volume ---
LV Path                /dev/vg01/lv01          1
LV Name                lv01
VG Name                vg01                          2
LV UUID                FVmNe1-u25R-dt3p-C5L6-VP2w-QRNP-scqrbq
LV Write Access        read/write
LV Creation host, time servera.lab.example.com, 2022-04-07 10:45:34 -0400
LV Status              available
# open                 1
LV Size                300.00 MiB          3
Current LE             75          4
Segments              1
Allocation             inherit
Read ahead sectors     auto
- currently set to    8192
Block device          253:0
```

- 1 LV Path shows the device name of the LV.
- 2 VG Name shows the VG that is used for creating this LV.
- 3 LV Size shows the total size of the LV. Use the file-system tools to determine the free and used space for the LV.
- 4 Current LE shows the number of logical extents that this LV uses.

Extend and Reduce LVM Storage

After creating a logical volume, you can extend the volume to expand the file system. You might need to extend PV or VG to increase the storage capacity of the LV.

Extend a Volume Group Size

You might need to add more disk space to extend a VG. You can add physical volumes to a VG to extend its available size.

Prepare the physical device and create the physical volume when not present.

```
[root@host ~]# parted /dev/vdb mkpart primary 1072MiB 1648MiB
...output omitted...
[root@host ~]# parted /dev/vdb set 3 lvm on
...output omitted...
[root@host ~]# udevadm settle
[root@host ~]# pvcreate /dev/vdb3
Physical volume "/dev/vdb3" successfully created.
```

The `vgextend` command adds the new PV to the VG. Provide the VG and PV names as arguments to the `vgextend` command.

```
[root@host ~]# vgextend vg01 /dev/vdb3
Volume group "vg01" successfully extended
```

This command extends the `vg01` VG by the `/dev/vdb3` PV size.

Extend a Logical Volume Size

A benefit of using logical volumes is increasing their size without experiencing any downtime. Add free physical extents to the LV from the VG, to extend the LV capacity and to expand the file system. Use the `vgdisplay` command to confirm that the volume group has enough free space for the LV extension.

Use the `lvextend` command to extend the LV.

```
[root@host ~]# lvextend -L +500M /dev/vg01/lv01
Size of logical volume vg01/lv01 changed from 300.00 MiB (75 extents) to 800.00
MiB (200 extents).
Logical volume vg01/lv01 successfully resized.
```

This command increases the size of the `lv01` logical volume by 500 MiB. The (+) plus sign in front of the size means adding this value to the existing size; otherwise, without the plus sign, the value defines the final size of the LV.

The `lvextend` command `-l` option expects the number of PE as the argument. The `lvextend` command `-L` option expects sizes in bytes, mebibytes, gibibytes, and similar.

Extend an XFS File System to the Logical Volume Size

The `xfs_growfs` command helps to expand the file system to occupy the extended LV. The target file system must be mounted before you use the `xfs_growfs` command. You can continue to use the file system when resizing.

```
[root@host ~]# xfs_growfs /mnt/data/
...output omitted...
data blocks changed from 76800 to 204800
```

**Important**

Always run the `xfs_growfs` command after executing the `lvextend` command. Use the `lvextend` command `-r` option to run the two steps consecutively. After extending the LV, resize the file system by using the `fsadm` command. This option supports several other file systems.

Extend an EXT4 File System to the Logical Volume Size

The steps for extending LV by using the `ext4` file system are the same as for LV by using the XFS file system, except for the step to resize the file system.

The `resize2fs` command expands the file system to occupy the new extended LV. You can continue to use the file system when resizing.

```
[root@host ~]# resize2fs /dev/vg01/lv01
resize2fs 1.46.5 (30-Dec-2021)
Resizing the filesystem on /dev/vg01/lv01 to 256000 (4k) blocks.
The filesystem on /dev/vg01/lv01 is now 256000 (4k) blocks long.
```

The `xfs_growfs` command expands an XFS file system. This command typically takes the mount point as an argument, or otherwise can indirectly use a device name by mounting the device and then specifying the mount point. In contrast, the `resize2fs` command resizes `ext2`, `ext3`, or `ext4` file systems and takes the device name as an argument. The `xfs_growfs` command supports only online resizing, whereas the `resize2fs` command supports both online and offline resizing. Although you can resize an `ext4` file system up or down, you can resize an XFS file system only up.

Extend Swap Space Logical Volumes

You can extend the LVs that are used as swap space; however, the process differs from expanding the `ext4` or XFS file system. Logical volumes that are used as swap space must be offline to extend them.

Use the `swaponoff` command to deactivate the swap space on the LV.

```
[root@host ~]# swaponoff -v /dev/vg01/swap
swaponoff /dev/vg01/swap
```

Use the `lvextend` command to extend the LV.

```
[root@host ~]# lvextend -L +300M /dev/vg01/swap
Size of logical volume vg01/swap changed from 500.00 MiB (125 extents) to 800.00
MiB (200 extents).
Logical volume vg01/swap successfully resized.
```

Use the `mkswap` command to format the LV as a swap space.

```
[root@host ~]# mkswap /dev/vg01/swap
mkswap: /dev/vg01/swap: warning: wiping old swap signature.
Setting up swapspace version 1, size = 800 MiB (838856704 bytes)
no label, UUID=25b4d602-6180-4b1c-974e-7f40634ad660
```

Use the `swapon` command to activate the swap space on the LV.

```
[root@host ~]# swapon /dev/vg01/swap
```

Reduce Volume Group Storage

Reducing a VG involves removing unused PVs from the VG. The `pvmove` command moves data from extents on one PV to extents on another PV with enough free extents in the same VG. You can continue to use the LV from the same VG when reducing. Use the `pvmove` command `-A` option to automatically back up the metadata of the VG after a change. This option uses the `vgcfgbackup` command to back up metadata automatically.

```
[root@host ~]# pvmove -A y /dev/vdb3
```



Warning

Before using the `pvmove` command, back up the data that is stored on all LVs in the VG. An unexpected power loss during the operation might leave the VG in an inconsistent state, which might cause a loss of data on LVs.

Use the `vgreduce` command to remove a PV from a VG.

```
[root@host ~]# vgreduce vg01 /dev/vdb3
Removed "/dev/vdb3" from volume group "vg01"
```



Important

The GFS2 and XFS file systems do not support shrinking, so you cannot reduce the size of an LV.

Remove LVM Storage

Use the `lvremove`, `vgremove`, and `pvremove` commands to remove an LVM component that is no longer required.

Prepare the File System

Move to another file system all data that must be kept. Use the `umount` command to unmount the file system, and then remove any `/etc/fstab` entries that are associated with this file system.

```
[root@host ~]# umount /mnt/data
```



Warning

Removing a logical volume destroys any data that is stored on the logical volume. Back up or move your data **before** you remove the logical volume.

Remove the Logical Volume

Use the `lvremove` *DEVICE-NAME* command to remove a logical volume that is no longer needed. Unmount the LV file system before running this command. The command prompts for confirmation before removing the LV.

```
[root@host ~]# lvremove /dev/vg01/lv01
Do you really want to remove active logical volume vg01/lv01? [y/n]: y
Logical volume "lv01" successfully removed.
```

The LV's physical extents are freed and available to assign to existing or new LVs in the volume group.

Remove the Volume Group

Use the `vgremove` *VG-NAME* command to remove a volume group that is no longer needed.

```
[root@host ~]# vgrremove vg01
Volume group "vg01" successfully removed
```

The VG's physical volumes are freed and available to assign to existing or new VGs on the system.

Remove the Physical Volumes

Use the `pvremove` command to remove physical volumes that are no longer needed. Use a space-delimited list of PV devices to remove more than one device at a time. This command deletes the PV metadata from the partition (or disk). The partition is now available for reallocation or reformatting.

```
[root@host ~]# pvremove /dev/vdb1 /dev/vdb2
Labels on physical volume "/dev/vdb1" successfully wiped.
Labels on physical volume "/dev/vdb2" successfully wiped.
```



References

`fdisk(8)`, `gdisk(8)`, `parted(8)`, `partprobe(8)`, `lvm(8)`, `pvcreate(8)`, `vgcreate(8)`, `lvcreate(8)`, `mkfs(8)`, `pvdisk(8)`, `vgdisplay(8)`, `lvdisplay(8)`, `vgextend(8)`, `lvextend(8)`, `xfs_growfs(8)`, `resize2fs(8)`, `swapoff(8)`, `mkswap(8)`, `swapon(8)`, `pvmove(8)`, `vgcfgbackup(8)`, `vgreduce(8)`, `lvremove(8)`, `vgrremove(8)`, and `pvremove(8)` man pages

For further information, refer to *Configuring and Managing Logical Volumes* at https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/configuring_and_managing_logical_volumes/index

► Guided Exercise

Create and Extend Logical Volumes

Create and extend a physical volume, volume group, logical volume, and an XFS file system. You also persistently mount the logical volume file system.

Outcomes

- Create physical volumes, volume groups, and logical volumes with LVM tools.
- Create file systems on logical volumes and persistently mount them.
- Extend the volume group to include an additional physical volume.
- Resize the logical volume when the file system is still mounted and in use.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start lvm-manage
```

Instructions

- 1. Log in to the `servera` machine as the `student` user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

- 2. Create the physical device partition on the `/dev/vdb` storage device.

- 2.1. Create two partitions of 256 MiB each and set to the Linux LVM type. Use the `first` and `second` names for these partitions.

```
[root@servera ~]# parted /dev/vdb mklabel gpt
Information: You may need to update /etc/fstab.
```

```
[root@servera ~]# parted /dev/vdb mkpart first 1MiB 258MiB
Information: You may need to update /etc/fstab.
```

```
[root@servera ~]# parted /dev/vdb set 1 lvm on
Information: You may need to update /etc/fstab.
```

```
[root@servera ~]# parted /dev/vdb mkpart second 258MiB 514MiB
Information: You may need to update /etc/fstab.
```

```
[root@servera ~]# parted /dev/vdb set 2 lvm on
Information: You may need to update /etc/fstab.
```

2.2. Register the new partitions with the kernel.

```
[root@servera ~]# udevadm settle
```

2.3. List the partitions on the `/dev/vdb` storage device. In the **Number** column, the 1 and 2 values correspond to the `/dev/vdb1` and `/dev/vdb2` device partitions. The **Flags** column indicates the partition type.

```
[root@servera ~]# parted /dev/vdb print
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 5369MB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name      Flags
  1      1049kB  271MB   269MB                   first     lvm
  2      271MB   539MB   268MB                   second    lvm
```

▶ 3. Label the two new partitions as physical volumes.

```
[root@servera ~]# pvcreate /dev/vdb1 /dev/vdb2
Physical volume "/dev/vdb1" successfully created.
Physical volume "/dev/vdb2" successfully created.
Creating devices file /etc/lvm/devices/system.devices
```

▶ 4. Create the `servera_group` volume group by using the two new PVs.

```
[root@servera ~]# vgcreate servera_group /dev/vdb1 /dev/vdb2
Volume group "servera_group" successfully created
```

▶ 5. Create the `servera_volume` logical volume with a size of 400 MiB. This command creates the `/dev/servera_group/servera_volume` LV without a file system.

```
[root@servera ~]# lvcreate -n servera_volume -L 400M servera_group
Logical volume "servera_volume" created.
```

▶ 6. Format the newly created LV and mount it persistently.

6.1. Format the `servera_volume` LV with the XFS file system.

```
[root@servera ~]# mkfs -t xfs /dev/servera_group/servera_volume
...output omitted...
```

- 6.2. Create the `/data` directory as a mount point.

```
[root@servera ~]# mkdir /data
```

- 6.3. To persistently mount the newly created file system, add the following content in the `/etc/fstab` file:

```
/dev/servera_group/servera_volume /data xfs defaults 0 0
```

- 6.4. Mount the `servera_volume` LV.

```
[root@servera ~]# mount /data
```

- 7. Verify that the mounted file system is accessible, and display the status information of the LVM.

- 7.1. Verify that you can copy files to the `/data` directory.

```
[root@servera ~]# cp -a /etc/*.conf /data
[root@servera ~]# ls /data | wc -l
32
```

- 7.2. View the PV status information. The output shows that the PV uses the `servera_group` VG. The PV has a size of 256 MiB and a physical extent size of 4 MiB.

The VG contains 63 PEs, of which 27 PEs are available for allocation, and 36 PEs are currently allocated to LVs. Use the following calculation for allocating the volume size in MiBs:

- Total 252 MiB (63 PEs x 4 MiB)
- Free 108 MiB (27 PEs x 4 MiB)
- Allocated 144 MiB (36 PEs x 4 MiB)

```
[root@servera ~]# pvdisplay /dev/vdb2
--- Physical volume ---
PV Name                /dev/vdb2
VG Name                servera_group
PV Size                256.00 MiB / not usable 4.00 MiB
Allocatable            yes
PE Size                4.00 MiB
Total PE               63
Free PE                27
Allocated PE           36
PV UUID                FKKFYJ-WJiR-1jt2-sfy3-yjPy-Tyln-LG92jj
```

- 7.3. View the VG status information of the `servera_group` VG. The output shows a VG size of 508 MiB with a PE size of 4 MiB. The available size from the VG is 108 MiB.

```
[root@servera ~]# vgdisplay servera_group
--- Volume group ---
VG Name                servera_group
```

```

System ID
Format                lvm2
Metadata Areas        2
Metadata Sequence No  2
VG Access              read/write
VG Status              resizable
MAX LV                0
Cur LV                1
Open LV                1
Max PV                0
Cur PV                2
Act PV                2
VG Size                508.00 MiB
PE Size                4.00 MiB
Total PE              127
Alloc PE / Size        100 / 400.00 MiB
Free PE / Size         27 / 108.00 MiB
VG UUID                g0ahyT-90J5-iGic-nnb5-G6T9-tLdK-dX8c9M

```

- 7.4. View the status information for the `servera_volume` LV. The output shows the VG name for creating the LV. It also shows an LV size of 400 MiB and an LE size of 100.

```

[root@servera ~]# lvmdisplay /dev/servera_group/servera_volume
--- Logical volume ---
LV Path                /dev/servera_group/servera_volume
LV Name                servera_volume
VG Name                servera_group
LV UUID                93MfUt-esgT-B5HM-r1p5-DVZH-n5cn-J5e2tw
LV Write Access        read/write
LV Creation host, time servera.lab.example.com, 2022-04-11 03:25:12 -0400
LV Status              available
# open                 1
LV Size                400.00 MiB
Current LE             100
Segments               2
Allocation             inherit
Read ahead sectors     auto
- currently set to    8192
Block device           253:0

```

- 7.5. View the free disk space in human-readable units. The output shows the total size of 395 MiB with the available size of 372 MiB.

```

[root@servera ~]# df -h /data
Filesystem                                Size  Used Avail Use% Mounted on
/dev/mapper/servera_group-servera_volume 395M   24M  372M   6% /data

```

- 8. Create the physical resource on the `/dev/vdb` storage device.

- 8.1. Create an additional partition of 512 MiB and set it to the Linux LVM type. Use the `third` name for this partition.


```
[root@servera ~]# parted /dev/vdb mkpart third 514MiB 1026MiB
[root@servera ~]# parted /dev/vdb set 3 lvm on
```

8.2. Register the new partition with the kernel.

```
[root@servera ~]# udevadm settle
```

8.3. Add the new partition as a PV.

```
[root@servera ~]# pvcreate /dev/vdb3
Physical volume "/dev/vdb3" successfully created.
```

- ▶ 9. Using the newly created disk space, extend the file system on the `servera_volume` to be a total size of 700 MiB.

9.1. Extend the `servera_group` VG by using the new `/dev/vdb3` PV.

```
[root@servera ~]# vgextend servera_group /dev/vdb3
Volume group "servera_group" successfully extended
```

9.2. Extend the existing `servera_volume` LV to 700 MiB.

```
[root@servera ~]# lvextend -L 700M /dev/servera_group/servera_volume
Size of logical volume servera_group/servera_volume changed from 400.00 MiB (100
extents) to 700.00 MiB (175 extents).
Logical volume servera_group/servera_volume successfully resized.
```

9.3. Extend the XFS file system by using the free space on the LV.

```
[root@servera ~]# xfs_growfs /data
...output omitted...
data blocks changed from 102400 to 179200
```

- ▶ 10. Verify that the LV size is extended, and that the contents are still present in the volume.

10.1. Verify the size of the extended LV by using the `lvdisplay` command.

```
[root@servera ~]# lvdisplay /dev/servera_group/servera_volume
--- Logical volume ---
LV Path                /dev/servera_group/servera_volume
LV Name                 servera_volume
VG Name                 servera_group
LV UUID                 mLQhsD-hyL0-KC2B-2nug-o2Nc-0znS-Q428fK
LV Write Access         read/write
LV Creation host, time  servera.lab.example.com, 2022-04-12 06:04:12 -0400
LV Status                available
# open                  1
LV Size                 700.00 MiB
Current LE               175
Segments                3
```

```
Allocation          inherit
Read ahead sectors  auto
- currently set to  8192
Block device        253:0
```

10.2. Verify the new file-system size. Verify that the previously copied files are still present.

```
[root@servera ~]# df -h /data
Filesystem                                Size  Used Avail Use% Mounted on
/dev/mapper/servera_group-servera_volume  695M   26M  670M   4% /data
[root@servera ~]# ls /data | wc -l
32
```

- 11. Return to the workstation machine as the student user.

```
[root@servera ~]# exit
logout
[student@servera ~]$ exit
logout
Connection to servera closed.
```

Finish

On the workstation machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish lvm-manage
```

Manage Layered Storage

Objectives

- Analyze the multiple storage components that make up the layers of the storage stack.

Storage Stack

Storage in RHEL is composed of multiple layers of drivers, managers, and utilities that are mature, stable, and full of modern features. Managing storage requires familiarity with stack components, and recognizing that storage configuration affects the boot process, application performance, and the ability to provide needed storage features for specific application use cases.




Figure 8.2: Storage stack

Previous sections in the Red Hat System Administration courses presented XFS file systems, network storage sharing, partitioning, and the Logical Volume Manager. This section shows the bottom-to-top RHEL storage stack and introduces each layer.

This section also covers Stratis, the daemon that unifies, configures, and monitors the underlying RHEL storage stack components, and provides automated local storage management from either the CLI or the RHEL web console.

Block Device

Block devices are at the bottom of the storage stack, and present a stable, consistent device protocol that enables including almost any block device transparently in a RHEL storage configuration. Most block devices today are accessed through the RHEL SCSI device driver, and appear as a SCSI device, including earlier ATA hard drives, solid-state devices, and common enterprise host bus adapters (HBAs). RHEL also supports iSCSI, Fibre Channel over Ethernet (FCoE), virtual machine driver (`virtio`), serial-attached SCSI (SAS), Non-Volatile Memory Express (NVMe), and other block devices.

An iSCSI target can be a dedicated physical device in a network or an iSCSI software-configured logical device on a networked storage server. The target is the endpoint in a SCSI protocol bus communication, to access the storage as Logical Unit Numbers (LUNs).

The Fibre Channel over Ethernet (FCoE) protocol transmits Fibre Channel frames over Ethernet networks. Typically, each data center has dedicated LAN and Storage Area Network (SAN) cabling, which is uniquely configured for its traffic. With FCoE, both traffic types can be combined

into a larger, converged, Ethernet network architecture. FCoE benefits include lower hardware and energy costs.

Multipath

A path is a connection between a server and the underlying storage. Device Mapper multipath (`dm-multipath`) is a RHEL native multipath tool for configuring redundant I/O paths into a single, path-aggregated logical device. A logical device that is created by using the device mapper (`dm`) appears as a unique block device in the `/dev/mapper/` directory for each LUN that is attached to the system.

You can also implement storage multipath redundancy by using network bonding when the storage, such as iSCSI and FCoE, uses network cabling.

Partitions

A block device can be further divided into partitions. Partitions might consume the entire block device size, or divide the block device for creating multiple partitions. You can use these partitions to create a file system, LVM devices, or directly for database structures or other raw storage.

RAID

A Redundant Array of Inexpensive Disks (RAID) is a storage virtualization technology that creates large logical volumes from multiple physical or virtual block device components. Different forms of RAID volumes offer data redundancy, performance improvement, or both, by implementing mirroring or striping layouts.

LVM supports RAID levels 0, 1, 4, 5, 6, and 10. RAID logical volumes that LVM creates and manages use the Multiple Devices (`mdadm`) kernel drivers. When not using LVM, Device Mapper RAID (`dm-raid`) provides a device mapper interface to the `mdadm` kernel drivers.

Logical Volume Manager

LVM physical volumes, volume groups, and logical volumes were discussed in a previous section. LVM can take almost any form of physical or virtual block devices, and can build storage as new logical storage volumes, and effectively hides the physical storage configuration from applications and other storage clients.

You can stack LVM volumes and implement advanced features such as encryption and compression for each part of the stack. The stack LVM volumes have mandated rules and recommended practices to follow for practical layering for specific scenarios. You can use case-specific recommendations from the *Configuring and Managing Logical Volumes* user guide.

LVM supports *LUKS encryption*, where a lower block device or partition is encrypted and presented as a secure volume to create a file system on top. The practical advantage for LUKS over file-system or file-based encryption is that a LUKS-encrypted device does not allow public visibility or access to the file-system structure. The LUKS-encrypted device ensures that a physical device remains secure even when removed from a computer.

LVM now incorporates VDO deduplication and compression as a configurable feature of regular logical volumes. You can use LUKS encryption and VDO together with logical volumes, where the LVM LUKS encryption is enabled underneath the LVM VDO volume.

File System or Other Use

The top layer of the stack is typically a file system, and can be used as raw space for databases or custom application data requirements. RHEL supports multiple file-system types, and

recommends XFS for most modern use cases. XFS is required when the utility that implements LVM is Red Hat Ceph Storage or the Stratis storage tool.

Database server applications consume storage in different ways, depending on their architecture and size. Some smaller databases store their structures in regular files that are contained in a file system. Because of the additional overhead or restrictions of file system access, this architecture has scaling limits. Larger databases that bypass file system caching, and that use their own caching mechanisms, create their database structures on raw storage. Logical volumes are suitable for database and other raw storage use cases.

Red Hat Ceph Storage creates its own storage management metadata structures on raw devices, to create Ceph Object Storage Devices (OSDs). In the latest Red Hat Ceph Storage versions, Ceph uses LVM to initialize disk devices for use as OSDs. More information is available in the *Cloud Storage with Red Hat Ceph Storage* (CL260) course.

Stratis Storage Management

Stratis is a local storage management tool that Red Hat and the upstream Fedora community developed. Stratis configures initial storage, changes a storage configuration, and uses advanced storage features.



Important

Stratis is currently available as a Technology Preview, but is expected to be supported in a later RHEL 9 version. For information about Red Hat scope of support for Technology Preview features, see the Technology Features Support Scope [<https://access.redhat.com/support/offerings/techpreview>] document.

Red Hat encourages customers to provide feedback when deploying Stratis.

Stratis runs as a service that manages pools of physical storage devices, and transparently creates and manages volumes for the newly created file systems.

Stratis builds file systems from shared pools of disk devices by using the *thin provisioning* concept. Instead of immediately allocating physical storage space to the file system when you create it, Stratis dynamically allocates that space from the pool as the file system stores more data. Therefore, the file system might appear to be 1 TiB, but might have only 100 GiB of real storage that is allocated to it from the pool.

You can create multiple pools from different storage devices. From each pool, you can create one or more file systems. Currently, you can create up to 2²⁴ file systems per pool.

Stratis builds the components that make up a Stratis-managed file system from standard Linux components. Internally, Stratis uses the Device Mapper infrastructure that LVM also uses. Stratis formats the managed file systems with XFS.

Figure 8.3 illustrates how Stratis assembles the elements of its storage management solution. Stratis assigns block storage devices such as hard disks or SSDs to pools. Each device contributes some physical storage to the pool. Then, Stratis creates file systems from the pools, and maps physical storage to each file system as needed.




Figure 8.3: Stratis architecture

Stratis Administration Methods

To manage file systems with the Stratis storage management solution, install the `stratis-cli` and `stratisd` packages. The `stratis-cli` package provides the `stratis` command, which sends reconfiguration requests to the `stratisd` system daemon. The `stratisd` package provides the `stratisd` service, which handles reconfiguration requests, and manages and monitors Stratis block devices, pools, and file systems.

Stratis administration is included in the RHEL web console.



Warning

Reconfigure file systems created by Stratis only with Stratis tools and commands.

Stratis uses stored metadata to recognize managed pools, volumes, and file systems. Manually configuring Stratis file systems with non-Stratis commands can result in overwriting that metadata, and can prevent Stratis from recognizing the file system volumes that it previously created.

Install and Enable Stratis

To use Stratis, ensure that your system has the required software and that the `stratisd` service is running. Install the `stratis-cli` and `stratisd` packages, and start and enable the `stratisd` service.

```
[root@host ~]# dnf install stratis-cli stratisd
...output omitted...
Is this ok [y/N]: y
...output omitted...
Complete!
[root@host ~]# systemctl enable --now stratisd
```

Create Stratis Pools

Create pools of one or more block devices by using the `stratis pool create` command. Then, use the `stratis pool list` command to view the list of available pools.

```
[root@host ~]# stratis pool create pool1 /dev/vdb
[root@host ~]# stratis pool list
```

Name	Total	Physical	Properties	UUID
pool1	5 GiB / 37.63 MiB / 4.96 GiB		~Ca,~Cr	11f6f3c5-5...

**Warning**

The `stratis pool list` command displays the storage space in use and the available pool space. Currently, if a pool becomes full, then further data that is written to the pool's file systems is quietly discarded.

Use the `stratis pool add-data` command to add block devices to a pool. Then, use the `stratis blockdev list` command to verify the block devices of a pool.

```
[root@host ~]# stratis pool add-data pool1 /dev/vdc
[root@host ~]# stratis blockdev list pool1
```

Pool Name	Device Node	Physical Size	Tier
pool1	/dev/vdb	5 GiB	Data
pool1	/dev/vdc	5 GiB	Data

Manage Stratis File Systems

Use the `stratis filesystem create` command to create a file system from a pool. The links to the Stratis file systems are in the `/dev/stratis/pool1` directory. Use the `stratis filesystem list` command to view the list of available file systems.

```
[root@host ~]# stratis filesystem create pool1 fs1
[root@host ~]# stratis filesystem list
```

Pool Name	Name	Used	Created	Device	UUID
pool1	fs1	546 MiB	Apr 08 2022 04:05	/dev/stratis/pool1/fs1	c7b5719...

Create a Stratis file system snapshot by using the `stratis filesystem snapshot` command. Snapshots are independent of the source file systems. Stratis dynamically allocates the snapshot storage space, and uses an initial 560 MB to store the file system's journal.

```
[root@host ~]# stratis filesystem snapshot pool1 fs1 snapshot1
```

Persistently Mount Stratis File Systems

You can persistently mount Stratis file systems by editing the `/etc/fstab` file and specifying the details of the file system. Use the `lsblk` command to display the UUID of the file system to use in the `/etc/fstab` file to identify the file system. You can also use the `stratis filesystem list` command to obtain the UUID of the file system.

```
[root@host ~]# lsblk --output=UUID /dev/stratis/pool1/fs1
```

UUID
c7b57190-8fba-463e-8ec8-29c80703d45e

The following example shows an entry in the `/etc/fstab` file to mount a Stratis file system persistently. This example entry is a single long line in the file. The `x-systemd.requires=stratisd.service` mount option delays mounting the file system until the `systemd` daemon starts the `stratisd` service during the boot process.

```
UUID=c7b57190-8fba-463e-8ec8-29c80703d45e /dir1 xfs defaults,x-
systemd.requires=stratisd.service 0 0
```



Important

If you do not include the `x-systemd.requires=stratisd.service` mount option in the `/etc/fstab` file for each Stratis file system, then the machine fails to start correctly, and aborts to `emergency.target` the next time that you reboot it.



Warning

Do not use the `df` command to query Stratis file system space.

The `df` command reports that any mounted Stratis-managed XFS file system is 1 TiB, regardless of the current allocation. Because the file system is thinly provisioned, a pool might not have enough physical storage to back the entire file system. Other file systems in the pool might use up all the available storage.

Therefore, it is possible to consume the whole storage pool, even if the `df` command reports that the file system has available space. Writes to a file system with no available pool storage can fail.

Instead, always use the `stratis pool list` command to monitor a pool's available storage accurately.



References

For further information, refer to *Deduplicating and Compressing Logical Volumes* on Red Hat Enterprise Linux at https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/deduplicating_and_compressing_logical_volumes_on_rhel/index

For further information, refer to *Red Hat Enterprise Linux 9 Managing File Systems Guide* at https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/managing_file_systems

Stratis Storage

<https://stratis-storage.github.io/>

What Stratis Learned from ZFS, Btrfs, and Linux Volume Manager

<https://opensource.com/article/18/4/stratis-lessons-learned>

► Guided Exercise

Manage Layered Storage

Use Stratis to create file systems from pools of storage that physical storage devices provide.

Outcomes

- Create a thin-provisioned file system by using the Stratis storage management solution.
- Verify that the Stratis volumes grow dynamically to support real-time data growth.
- Access data from the snapshot of a thin-provisioned file system.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start lvm-stratis
```

Instructions

- 1. Log in to the servera machine as the student user and switch to the root user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

- 2. Install the `stratisd` and `stratis-cli` packages.

```
[root@servera ~]# dnf install stratisd stratis-cli
...output omitted...
Is this ok [y/N]: y
...output omitted...
Complete!
```

- 3. Activate the `stratisd` service.

```
[root@servera ~]# systemctl enable --now stratisd
```

- 4. Ensure that the `stratispool1` Stratis pool exists on the `/dev/vdb` block device.

- 4.1. Create the `stratispool1` Stratis pool.

```
[root@servera ~]# stratis pool create stratispool1 /dev/vdb
```

- 4.2. Verify the availability of the `stratispool1` pool. Note the size of the pool.

```
[root@servera ~]# stratis pool list
```

Name	Total Physical	Properties	UUID
stratispool1	5 GiB / 37.63 MiB / 4.96 GiB	~Ca, ~Cr	3557c389-7...

- ▶ 5. Expand the capacity of the `stratispool1` pool by adding the `/dev/vdc` block device.

- 5.1. Add the `/dev/vdc` block device to the `stratispool1` pool.

```
[root@servera ~]# stratis pool add-data stratispool1 /dev/vdc
```

- 5.2. Verify the size of the `stratispool1` pool. The `stratispool1` pool size increases when you add the block device.

```
[root@servera ~]# stratis pool list
```

Name	Total Physical	Properties	UUID
stratispool1	10 GiB / 41.63 MiB / 9.96 GiB	~Ca, ~Cr	3557c389-7...

- 5.3. Verify the block devices that are currently members of the `stratispool1` pool.

```
[root@servera ~]# stratis blockdev list stratispool1
```

Pool Name	Device Node	Physical Size	Tier
stratispool1	/dev/vdb	5 GiB	Data
stratispool1	/dev/vdc	5 GiB	Data

- ▶ 6. Add a thin-provisioned file system called `stratis-filesystem1` in the `stratispool1` pool. Mount the file system on the `/stratisvol` directory. Create a file on the `stratis-filesystem1` file system called `file1` that contains the text `Hello World!`. Modify the `/etc/fstab` file to persistently mount the file system on the `/stratisvol` directory.

- 6.1. Create the thin-provisioned `stratis-filesystem1` file system on the `stratispool1` pool. It might take up to a minute for the command to complete.

```
[root@servera ~]# stratis filesystem create stratispool1 stratis-filesystem1
```

- 6.2. Verify the availability of the `stratis-filesystem1` file system, and note its current usage. The usage of the file system increases on demand in the later steps.

```
[root@servera ~]# stratis filesystem list
```

Pool Name	Name	Used	Created	Device
stratispool1	stratis-filesystem1	546 MiB	Apr 08 2022 07:12	/dev/stratis/
stratispool1/stratis-filesystem1		48e8...		

- 6.3. Create the `/stratisvol` directory.

```
[root@servera ~]# mkdir /stratisvol
```

6.4. Mount the `stratis-filesystem1` file system on the `/stratisvol` directory.

```
[root@servera ~]# mount /dev/stratis/stratispool1/stratis-filesystem1 \
/stratisvol
```

6.5. Create the `/stratisvol/file1` text file.

```
[root@servera ~]# echo "Hello World!" > /stratisvol/file1
```

6.6. Unmount the `/stratisvol` volume.

```
[root@servera ~]# umount /stratisvol
```

6.7. Obtain the UUID of the file system. The UUID would be different in your system.

```
[root@servera ~]# lsblk --output=UUID \
/dev/stratis/stratispool1/stratis-filesystem1
UUID
d18cb4fc-753c-473a-9ead-d6661533b475
```

6.8. Modify the `/etc/fstab` file to persistently mount the file system on the `/stratisvol` directory. To do so, use the `vim /etc/fstab` command and add the following line. Replace the UUID with the correct one for your system.

```
UUID=d18c... /stratisvol xfs defaults,x-systemd.requires=stratisd.service 0 0
```

6.9. Update the `systemd` daemon with the new `/etc/fstab` configuration file.

```
[root@servera ~]# systemctl daemon-reload
```

6.10. Mount the `stratisvol` volume and verify that the `stratis-filesystem1` volume is mounted on the `/stratisvol` directory.

```
[root@servera ~]# mount /stratisvol
[root@servera ~]# mount
...output omitted...
/dev/mapper/stratis-1-3557...fbd3-thin-fs-48e8...9ebe on /stratisvol type xfs
(rw,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,sunit=2048,swidth=2048,
noquota)
```

► 7. Reboot your system and verify that the file system is persistently mounted across reboots.

```
[root@servera ~]# systemctl reboot
...output omitted...
[student@workstation ~]$ ssh student@servera
```

```
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]# mount
...output omitted...
/dev/mapper/stratis-1-3557...fbd3-thin-fs-d18c...b475 on /stratisvol type xfs
(rw,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,sunit=2048,swidth=2048,
noquota,x-systemd.requires=stratisd.service)
```

- 8. Verify that the `stratis-filesystem1` thin-provisioned file system dynamically grows as the data on the file system grows.

8.1. View the current usage of the `stratis-filesystem1` file system.

```
[root@servera ~]# stratis filesystem list
```

Pool Name	Name	Used	Created	Device
		UUID		
stratispool1	stratis-filesystem1	546 MiB	Apr 08 2022 07:12	/dev/stratis/
stratispool1/stratis-filesystem1		48e8...		

- 8.2. Create a 2 GiB file on the `stratis-filesystem1` file system. It might take up to a minute for the command to complete.

```
[root@servera ~]# dd if=/dev/urandom of=/stratisvol/file2 bs=1M count=2048
```

- 8.3. Verify the used space in the `stratis-filesystem1` file system.

The output shows that the used space in the `stratis-filesystem1` file system increased. The used-space increase confirms that the thin-provisioned file system dynamically expands as needed.

```
[root@servera ~]# stratis filesystem list
```

Pool Name	Name	Used	Created	Device
		UUID		
stratispool1	stratis-filesystem1	2.60 GiB	Apr 08 2022 07:12	/dev/stratis/
stratispool1/stratis-filesystem1		48e8...		

- 9. Create a snapshot called `stratis-filesystem1-snap` of the `stratis-filesystem1` file system. The snapshot provides you with access to any file that you delete from the `stratis-filesystem1` file system.

- 9.1. Create a snapshot of the `stratis-filesystem1` file system. It might take up to a minute for the command to complete.

```
[root@servera ~]# stratis filesystem snapshot stratispool1 \
stratis-filesystem1 stratis-filesystem1-snap
```

- 9.2. Verify the availability of the snapshot.

```
[root@servera ~]# stratis filesystem list
```

Pool Name	Name	Used UUID	Created	Device
stratispool1	stratis-filesystem1-snap	2.73 GiB	Apr 08 2022 07:22	/dev/stratis/stratispool1/stratis-filesystem1-snap
stratispool1	stratis-filesystem1	2.73 GiB	Apr 08 2022 07:12	/dev/stratis/stratispool1/stratis-filesystem1

9.3. Remove the `/stratisvol/file1` file.

```
[root@servera ~]# rm /stratisvol/file1
rm: remove regular file '/stratisvol/file1'? y
```

9.4. Create the `/stratisvol-snap` directory.

```
[root@servera ~]# mkdir /stratisvol-snap
```

9.5. Mount the `stratis-filesystem1-snap` snapshot on the `/stratisvol-snap` directory.

```
[root@servera ~]# mount /dev/stratis/stratispool1/stratis-filesystem1-snap \
/stratisvol-snap
```

9.6. Verify that you can still access the file that you deleted from the `stratis-filesystem1` file system in the snapshot.

```
[root@servera ~]# cat /stratisvol-snap/file1
Hello World!
```

► 10. Unmount the `/stratisvol` and `/stratisvol-snap` volumes.

```
[root@servera ~]# umount /stratisvol-snap
[root@servera ~]# umount /stratisvol
```

► 11. Remove the `stratis-filesystem1` thin-provisioned file system and the `stratis-filesystem1-snap` snapshot from the system.

11.1. Destroy the `stratis-filesystem1-snap` snapshot.

```
[root@servera ~]# stratis filesystem destroy stratispool1 stratis-filesystem1-snap
```

11.2. Destroy the `stratis-filesystem1` file system.

```
[root@servera ~]# stratis filesystem destroy stratispool1 stratis-filesystem1
```

11.3. Return to the `workstation` system as the `student` user.

```
[root@servera ~]# exit
logout
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish lvm-stratis
```

► Lab

Manage Storage Stack

Resize an existing logical volume, add LVM resources as necessary, and then add a new logical volume with a persistently mounted XFS file system on it.

Outcomes

- Resize the `serverb_01_lv` logical volume to 768 MiB.
- Create the `serverb_02_lv` logical volume with 128 MiB with an XFS file system.
- Persistently mount the volume on the `/storage/data2` directory.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start lvm-review
```

Instructions

On the `serverb` machine, the `serverb_01_lv` logical volume that is mounted on the `/storage/data1` directory is running out of disk space, and must be extended to 768 MiB. You must ensure that the `serverb_01_lv` LV remains persistently mounted on the `/storage/data1` directory.

The `serverb_01_lv` LV is present on the `serverb_01_vg` volume group. It has insufficient space to extend the existing logical volume. A 512 MiB partition exists on the `/dev/vdb` disk. Create a partition with the 512 MiB size on the `/dev/vdb` disk.



Important

Note especially the specification of the partition size in MiB (2^{20} bytes). If you create the partition in MB (10^6 bytes), it does not satisfy the evaluation criteria, because 1 MiB = 1.048576 MB.

Although the default unit when using the `parted /dev/vdb print` command is MB, you can verify the size of the `/dev/vdb` device partitions in MiB units. Use the `parted /dev/vdb unit MiB print` command to print the partition sizes in MiB.

Create the `serverb_02_lv` LV with 128 MiB. Create the XFS file system on the newly created volume. Mount the newly created logical volume on the `/storage/data2` directory.

1. Create a 512 MiB partition on the `/dev/vdb` disk. Initialize this partition as a physical volume, and extend the `serverb_01_vg` volume group to use this partition.

2. Extend the `serverb_01_lv` logical volume to 768 MiB.
3. In the existing volume group, create the `serverb_02_lv` logical volume with 128 MiB. Add an XFS file system and mount it persistently on the `/storage/data2` directory.
4. Verify that the newly created LV is mounted with the intended size.

Evaluation

As the student user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade lvm-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish lvm-review
```


► Solution

Manage Storage Stack

Resize an existing logical volume, add LVM resources as necessary, and then add a new logical volume with a persistently mounted XFS file system on it.

Outcomes

- Resize the `serverb_01_lv` logical volume to 768 MiB.
- Create the `serverb_02_lv` logical volume with 128 MiB with an XFS file system.
- Persistently mount the volume on the `/storage/data2` directory.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start lvm-review
```

Instructions

On the `serverb` machine, the `serverb_01_lv` logical volume that is mounted on the `/storage/data1` directory is running out of disk space, and must be extended to 768 MiB. You must ensure that the `serverb_01_lv` LV remains persistently mounted on the `/storage/data1` directory.

The `serverb_01_lv` LV is present on the `serverb_01_vg` volume group. It has insufficient space to extend the existing logical volume. A 512 MiB partition exists on the `/dev/vdb` disk. Create a partition with the 512 MiB size on the `/dev/vdb` disk.



Important

Note especially the specification of the partition size in MiB (2^{20} bytes). If you create the partition in MB (10^6 bytes), it does not satisfy the evaluation criteria, because 1 MiB = 1.048576 MB.

Although the default unit when using the `parted /dev/vdb print` command is MB, you can verify the size of the `/dev/vdb` device partitions in MiB units. Use the `parted /dev/vdb unit MiB print` command to print the partition sizes in MiB.

Create the `serverb_02_lv` LV with 128 MiB. Create the XFS file system on the newly created volume. Mount the newly created logical volume on the `/storage/data2` directory.

1. Create a 512 MiB partition on the `/dev/vdb` disk. Initialize this partition as a physical volume, and extend the `serverb_01_vg` volume group to use this partition.

- 1.1. Log in to the `serverb` machine as the `student` user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
[student@serverb ~]$ sudo -i
[sudo] password for student: student
[root@serverb ~]#
```

- 1.2. Print the partition sizes in MiB to determine where the first partition ends.

```
[root@serverb ~]# parted /dev/vdb unit MiB print
...output omitted...

Number  Start   End     Size    File system  Name      Flags
  1      1.00MiB 513MiB  512MiB                primary
```

- 1.3. Create the 512 MiB partition and set the `lvm` partition type.

```
[root@serverb ~]# parted /dev/vdb mkpart primary 514MiB 1026MiB
...output omitted...
[root@serverb ~]# parted /dev/vdb set 2 lvm on
```

- 1.4. Register the new partition with the kernel.

```
[root@serverb ~]# udevadm settle
```

- 1.5. Initialize the partition as a PV.

```
[root@serverb ~]# pvcreate /dev/vdb2
Physical volume "/dev/vdb2" successfully created.
```

- 1.6. Extend the `serverb_01_vg` VG by using the new `/dev/vdb2` PV.

```
[root@serverb ~]# vgextend serverb_01_vg /dev/vdb2
Volume group "serverb_01_vg" successfully extended
```

2. Extend the `serverb_01_lv` logical volume to 768 MiB.

- 2.1. Extend the `serverb_01_lv` LV to 768 MiB.

Alternatively, you can also use the `lvcreate` command `-L +512M` option to resize the LV.

```
[root@serverb ~]# lvextend -L 768M /dev/serverb_01_vg/serverb_01_lv
Size of logical volume serverb_01_vg/serverb_01_lv changed from 256.00 MiB (64
extents) to 768.00 MiB (192 extents).
Logical volume serverb_01_vg/serverb_01_lv successfully resized.
```

- 2.2. Extend the XFS file system to consume the remaining space on the LV.

```
[root@serverb ~]# xfs_growfs /storage/data1
meta-data=/dev/mapper/serverb_01_vg-serverb_01_lv isize=512    agcount=4,
          agsize=16384 blks
...output omitted...
data blocks changed from 65536 to 196608
```

**Note**

The `xfs_growfs` command introduces an extra step to extend the file system. An alternative would be to use the `lvextend` command `-r` option.

3. In the existing volume group, create the `serverb_02_lv` logical volume with 128 MiB. Add an XFS file system and mount it persistently on the `/storage/data2` directory.

- 3.1. Create the `serverb_02_lv` LV with 128 MiB from the `serverb_01_vg` VG.

```
[root@serverb ~]# lvcreate -n serverb_02_lv -L 128M serverb_01_vg
Logical volume "serverb_02_lv" created.
```

- 3.2. Create the xfs file system on the `serverb_02_lv` LV.

```
[root@serverb ~]# mkfs -t xfs /dev/serverb_01_vg/serverb_02_lv
...output omitted...
```

- 3.3. Create the `/storage/data2` directory as the mount point.

```
[root@serverb ~]# mkdir /storage/data2
```

- 3.4. Add the following line to the end of the `/etc/fstab` file:

```
/dev/serverb_01_vg/serverb_02_lv /storage/data2 xfs defaults 0 0
```

- 3.5. Update the `systemd` daemon with the new `/etc/fstab` configuration file.

```
[root@serverb ~]# systemctl daemon-reload
```

- 3.6. Mount the `serverb_02_lv` LV.

```
[root@serverb ~]# mount /storage/data2
```

4. Verify that the newly created LV is mounted with the intended size.

- 4.1. Use the `df` command to verify the `serverb_01_lv` LV size.

```
[root@serverb ~]# df -h /storage/data1
Filesystem                                Size  Used Avail Use% Mounted on
/dev/mapper/serverb_01_vg-serverb_01_lv  763M   19M  744M   3% /storage/data1
```

4.2. Verify the `serverb_02_lv` LV size.

```
[root@serverb ~]# df -h /storage/data2
Filesystem                                Size  Used Avail Use% Mounted on
/dev/mapper/serverb_01_vg-serverb_02_lv  123M   7.6M  116M   7% /storage/data2
```

4.3. Verify the `serverb_01_lv` LV details.

```
[root@serverb ~]# lvsdisplay /dev/serverb_01_vg/serverb_01_lv
--- Logical volume ---
LV Path                /dev/serverb_01_vg/serverb_01_lv
LV Name                 serverb_01_lv
VG Name                 serverb_01_vg
LV UUID                 1pY3DZ-fs1F-mptC-fL32-e8tG-PFBT-bs7LSJ
LV Write Access         read/write
LV Creation host, time serverb.lab.example.com, 2022-05-05 14:40:51 -0400
LV Status                available
# open                  1
LV Size                 768.00 MiB
Current LE              192
Segments                2
Allocation               inherit
Read ahead sectors      auto
- currently set to      8192
Block device            253:0
```

4.4. Verify the `serverb_02_lv` LV details.

```
[root@serverb ~]# lvsdisplay /dev/serverb_01_vg/serverb_02_lv
--- Logical volume ---
LV Path                /dev/serverb_01_vg/serverb_02_lv
LV Name                 serverb_02_lv
VG Name                 serverb_01_vg
LV UUID                 0aJIb6-Ti2b-jLCK-imB6-rkLx-mUoX-acjkz9
LV Write Access         read/write
LV Creation host, time serverb.lab.example.com, 2022-05-05 14:45:46 -0400
LV Status                available
# open                  1
LV Size                 128.00 MiB
Current LE              32
Segments                1
Allocation               inherit
Read ahead sectors      auto
- currently set to      8192
Block device            253:1
```

Evaluation

As the student user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade lvm-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish lvm-review
```

Summary

- You can use LVM to create flexible storage by allocating space on multiple storage devices.
- Physical volumes, volume groups, and logical volumes are managed by the `pvcreate`, `vgreduce`, and `lvextend` commands.
- Logical volumes can be formatted with a file system or swap space, and they can be mounted persistently.
- Storage can be added to volume groups, and logical volumes can be extended dynamically.
- Storage stack uses layers and components to manage storage efficiently.
- Virtual Data Optimizer (VDO) uses LVM for compression and deduplication of data.
- You can use Stratis to configure initial storage or to enable advanced storage features.

Chapter 9

Access Network-Attached Storage

Goal

Access network-attached storage with the NFS protocol.

Objectives

- Identify NFS export information, create a directory to use as a mount point, mount an NFS export with the `mount` command or by configuring the `/etc/fstab` file, and unmount an NFS export with the `umount` command.
- Describe the benefits of using the automounter, and automount NFS exports by using direct and indirect maps.

Sections

- Manage Network-Attached Storage with NFS (and Guided Exercise)
- Automount Network-Attached Storage (and Guided Exercise)

Lab

- Access Network-Attached Storage

Manage Network-Attached Storage with NFS

Objectives

- Identify NFS export information, create a directory to use as a mount point, mount an NFS export with the `mount` command or by configuring the `/etc/fstab` file, and unmount an NFS export with the `umount` command.

Accessing Exported NFS Directories

The *Network File System* (NFS) is an internet standard protocol that Linux, UNIX, and similar operating systems use as their native network file system. NFS is an open standard that supports native Linux permissions and file-system attributes.

By default, Red Hat Enterprise Linux 9 uses NFS version 4.2. RHEL fully supports both NFSv3 and NFSv4 protocols. NFSv3 might use either a TCP or a UDP transport protocol, but NFSv4 supports only TCP connections.

NFS servers *export* directories. NFS clients mount exported directories to an existing local mount point directory. NFS clients can mount exported directories in multiple ways:

- **Manually** by using the `mount` command.
- **Persistently at boot** by configuring entries in the `/etc/fstab` file.
- **On demand** by configuring an automounter method.

The automounter methods, which include the `autofs` service and the `systemd.automount` facility, are discussed in the **Automount Network-Attached Storage** section. You must install the `nfs-utils` package to obtain the client tools for manually mounting, or for automounting, to obtain exported NFS directories.

```
[root@host ~]# dnf install nfs-utils
```

RHEL also supports mounting *shared* directories from Microsoft Windows systems by using the same methods as for the NFS protocol, by using either the Server Message Block (SMB) or the Common Internet File System (CIFS) protocols. Mounting options are protocol-specific, and depend on your Windows Server or Samba Server configuration.

Query a Server's Exported NFS Directories

The NFS protocol changed significantly between NFSv3 and NFSv4. The method to query a server to view the available exports is different for each protocol version.

NFSv3 used the RPC protocol, which requires a file server that supports NFSv3 connections to run the `rpcbind` service. An NFSv3 client connects to the `rpcbind` service at port 111 on the server to request NFS service. The server responds with the current port for the NFS service. Use the `showmount` command to query the available exports on an RPC-based NFSv3 server.

```
[root@host ~]# showmount --exports server
```

```
Export list for server
/shares/test1
/shares/test2
```

The NFSv4 protocol eliminated the use of the legacy RPC protocol for NFS transactions. Use of the `showmount` command on a server that supports only NFSv4 times out without receiving a response, because the `rpcbind` service is not running on the server. However, querying an NFSv4 server is simpler than querying an NFSv3 server.

NFSv4 introduced an *export tree* that contains all of the paths for the server's exported directories. To view all of the exported directories, mount the root (/) of the server's export tree. Mounting the export tree's root provides browseable paths for all exported directories, as children of the tree's root directory, but does not mount ("bind") any of the exported directories.

```
[root@host ~]# mkdir /mountpoint
[root@host ~]# mount server:/ /mountpoint
[root@host ~]# ls /mountpoint
```

To mount an NFSv4 export when browsing the mounted export tree, change directory to an exported directory path. Alternatively, use the `mount` command with an exported directory's full path name to mount a single exported directory. Exported directories that use Kerberos security do not allow mounting or accessing a directory when browsing an export tree, even though you can view the export's path name. Mounting Kerberos-protected shares requires additional server configuration and the use of Kerberos user credentials, which are discussed in the *Red Hat Security: Identity Management and Active Directory Integration (RH362)* training course.

Manually Mount Exported NFS Directories

After you identify the NFS export to mount, create a local mount point if it does not yet exist. Although the `/mnt` directory is available for use as a temporary mount point, recommended practice is not to use `/mnt` for long-term or persistent mounting.

```
[root@host ~]# mkdir /mountpoint
```

As with local volume file systems, mount the NFS export to access its contents. NFS shares can be mounted temporarily or permanently, only by a privileged user.

```
[root@host ~]# mount -t nfs -o rw,sync server:/export /mountpoint
```

The `-t nfs` option specifies the NFS file-system type. However, when the `mount` command detects the `server:/export` syntax, the command defaults to the NFS type. With the `-o` flag, you can add a list of comma-separated options to the `mount` command. In the example, the `rw` option specifies that the exported file system is mounted with read/write access. The `sync` option specifies synchronous transactions to the exported file system. This method is strongly recommended for all production network mounts where transactions must be completed or else return as failed.

Using a manual `mount` command is not persistent. When the system reboots, that NFS export is not still mounted. Manual mounts are useful for providing temporary access to an exported directory, or for test mounting an NFS export before persistently mounting it.

Persistently Mount Exported NFS Directories

To persistently mount an NFS export, edit the `/etc/fstab` file, and add the mount entry with similar syntax to manual mounting.

```
[root@host ~]# vim /etc/fstab
...
server:/export /mountpoint nfs rw 0 0
```

Then, you can mount the NFS export by using only the mount point. The `mount` command obtains the NFS server and mount options from the matching entry in the `/etc/fstab` file.

```
[root@host ~]# mount /mountpoint
```

Unmount Exported NFS Directories

As a privileged user, unmount an NFS export with the `umount` command. Unmounting a share does not remove its entry in the `/etc/fstab` file, if that file exists. Entries in the `/etc/fstab` file are persistent and are remounted during boot.

```
[root@host ~]# umount /mountpoint
```

A mounted directory can sometimes fail to unmount, and returns an error that the device is busy. The device is busy because either an application is keeping a file open within the file system, or some user's shell has a working directory in the mounted file-system's root directory or below it.

To resolve the error, check your own active shell windows, and use the `cd` command to leave the mounted file system. If subsequent attempts to unmount the file system still fail, then use the `lsof` (*list open files*) command to query the mount point. The `lsof` command returns a list of open file names and the process which is keeping the file open.

```
[root@host ~]# lsof /mountpoint
COMMAND PID  USER FD  TYPE DEVICE SIZE/OFF  NODE NAME
program  5534  user  txt REG   252.4  910704    128 /home/user/program
```

With this information, gracefully close any processes that are using files on this file system, and retry the unmount. In critical scenarios only, when an application cannot be closed gracefully, kill the process to close the file. Alternatively, use the `umount -f` option to force the unmount, which can cause loss of unwritten data for all open files.



References

`mount(8)`, `umount(8)`, `showmount(8)`, `fstab(5)`, `mount.nfs(8)`, `nfsconf(8)`, and `rpcbind(8)` man pages

► Guided Exercise

Manage Network-Attached Storage with NFS

Modify the `/etc/fstab` file to persistently mount an NFS export at boot time.

Outcomes

- Test an NFS server with the `mount` command.
- Configure NFS exports in the `/etc/fstab` configuration file to save changes even after a system reboots.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start netstorage-nfs
```

Instructions

A shipping company uses a central NFS server, `serverb`, to host various exported documents and directories. Users on `servera`, who are all members of the `admin` group, need access to the persistently mounted NFS export.

The following list provides the environment characteristics for completing this exercise:

- The `serverb` machine exports the `/shares/public` directory, which contains some text files.
- Members of the `admin` group (`admin1`, `sysmanager1`) have read and write access to the `/shares/public` exported directory.
- The mount point on `servera` must be the `/public` directory.
- All user passwords are set to `redhat`.
- The `nfs-utils` package is already installed.

- 1. Log in to `servera` as the `student` user and switch to the `root` user.

- 1.1. Log in to `servera` as the `student` user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

► 2. Test the NFS server on `serverb` with `servera` as the NFS client.

2.1. Create the `/public` mount point on the `servera` machine.

```
[root@servera ~]# mkdir /public
```

2.2. On `servera`, verify that the `/shares/public` NFS export from `serverb` successfully mounts to the `/public` directory.

```
[root@servera ~]# mount -t nfs \
serverb.lab.example.com:/shares/public /public
```

2.3. List the contents of the mounted NFS export.

```
[root@servera ~]# ls -l /public
total 16
-rw-r--r--. 1 root admin 42 Apr  8 22:36 Delivered.txt
-rw-r--r--. 1 root admin 46 Apr  8 22:36 NOTES.txt
-rw-r--r--. 1 root admin 20 Apr  8 22:36 README.txt
-rw-r--r--. 1 root admin 27 Apr  8 22:36 Trackings.txt
```

2.4. Explore the `mount` command options for the mounted NFS export.

```
[root@servera ~]# mount | grep public
serverb.lab.example.com:/shares/public on /public type nfs4
(rw,relatime,vers=4.2,rsz=262144,wsz=262144,namlen=255,
sync,proto=tcp,timeo=600,retrans=2,sec=sys,clientaddr=172.25.250.10,
local_lock=none,addr=172.25.250.11)
```

2.5. Unmount the NFS export.

```
[root@servera ~]# umount /public
```

► 3. Configure `servera` so that the `/shares/public` export is persistently mounted.

3.1. Edit the `/etc/fstab` file.

```
[root@servera ~]# vim /etc/fstab
```

Add the following line to the end of the file:

```
serverb.lab.example.com:/shares/public /public nfs rw,sync 0 0
```

3.2. Mount the exported directory.

```
[root@servera ~]# mount /public
```

3.3. List the contents of the exported directory.

```
[root@servera ~]# ls -l /public
total 16
-rw-r--r--. 1 root    admin 42 Apr  8 22:36 Delivered.txt
-rw-r--r--. 1 root    admin 46 Apr  8 22:36 NOTES.txt
-rw-r--r--. 1 root    admin 20 Apr  8 22:36 README.txt
-rw-r--r--. 1 root    admin 27 Apr  8 22:36 Trackings.txt
```

3.4. Reboot the servera machine.

```
[root@servera ~]# systemctl reboot
```

- 4. After servera is finished rebooting, log in to servera as the admin1 user and test the persistently mounted NFS export.

4.1. Log in to servera as the admin1 user.

```
[student@workstation ~]$ ssh admin1@servera
[admin1@servera ~]$
```

4.2. Test the NFS export that is mounted on the /public directory.

```
[admin1@servera ~]$ ls -l /public
total 16
-rw-r--r--. 1 root    admin 42 Apr  8 22:36 Delivered.txt
-rw-r--r--. 1 root    admin 46 Apr  8 22:36 NOTES.txt
-rw-r--r--. 1 root    admin 20 Apr  8 22:36 README.txt
-rw-r--r--. 1 root    admin 27 Apr  8 22:36 Trackings.txt
[admin1@servera ~]$ cat /public/NOTES.txt
###In this file you can log all your notes###
[admin1@servera ~]$ echo "This is a test" > /public/Test.txt
[admin1@servera ~]$ cat /public/Test.txt
This is a test
```

4.3. Return to the workstation machine as the student user.

```
[admin1@servera ~]$ exit
logout
Connection to servera closed.
```

Finish

On the workstation machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish netstorage-nfs
```

Automount Network-Attached Storage

Objectives

- Describe the benefits of using the automounter, and automount NFS exports by using direct and indirect maps.

Mount NFS Exports with the Automounter

The *automounter* is a service (`autofs`) that automatically mounts file systems and NFS exports on demand, and automatically unmounts file systems and NFS exports when the mounted resources are no longer in current use.

The automounter function was created to solve the problem that unprivileged users do not have sufficient permissions to use the `mount` command. Without use of the `mount` command, normal users cannot access removable media such as CDs, DVDs, and removable disk drives. Furthermore, if a local or remote file system is not mounted at boot time by using the `/etc/fstab` configuration, then a normal user cannot mount and access those unmounted file systems.

The automounter configuration files are populated with file-system mount information, in a similar way to `/etc/fstab` entries. Although `/etc/fstab` file systems mount during system boot and remain mounted until system shutdown or other intervention, automounter file systems do not necessarily mount during system boot. Instead, automounter-controlled file systems mount on demand, when a user or application attempts to enter the file-system mount point to access files.

Automounter Benefits

Resource use for automounter file systems is equivalent to file systems that are mounted at boot, because a file system uses resources only when a program is reading and writing open files. Mounted but idle file systems and unmounted file systems use almost no resources.

The automounter advantage is that by unmounting the file system each time that it is no longer in use, the file system is protected from unexpected corruption when it is open. When the file system is directed to mount again, the `autofs` service uses the most current mount configuration, unlike an `/etc/fstab` mount, which might still use a configuration that was mounted months ago during the last system boot. Additionally, if your NFS server configuration includes redundant servers and paths, then the automounter can select the fastest connection each time that a new file system is requested.

The Automounter `autofs` Service Method

The `autofs` service supports the same local and remote file systems as in the `/etc/fstab` file, including NFS and SMB file sharing protocols, and supports the same protocol-specific mount options, including security parameters. File systems that are mounted through the automounter are available by default to all users, but can be restricted through access permission options.

Because the automounter is a client-side configuration that uses the standard `mount` and `umount` commands to manage file systems, automounted file systems in use exhibit the same behavior to file systems that are mounted by using `/etc/fstab`. The difference is that an automounter file system remains unmounted until the mount point is accessed, which causes the file system to mount immediately, and to remain mounted when the file system is in use. When all files on the file

system are closed, and all users and processes leave the mount point directory, the automounter unmounts the file system after a minimal timeout.

Direct and Indirect Map Use Cases

The automounter supports both direct and indirect mount-point mapping, to handle the two types of demand mounting. A *direct* mount is when a file system mounts to an unchanging, known mount point location. Almost all the file system mounts that you configured, before learning about the automounter, are examples of direct mounts. A *direct* mount point exists as a permanent directory, the same as other normal directories.

An *indirect* mount is when the mount point location is not known until the mount demand occurs. An example of an indirect mount is the configuration for remote-mounted home directories, where a user's home directory includes their username in the directory path. The user's remote file system is mounted to their home directory, only after the automounter learns which user specified to mount their home directory, and determines the mount point location to use. Although *indirect* mount points appear to exist, the `autofs` service creates them when the mount demand occurs, and deletes them again when the demand ended and the file system is unmounted.

Configure the Automounter Service

To use the `autofs` service, you must install the `autofs` and `nfs-utils` packages.

```
[root@host ~]# dnf install autofs nfs-utils
```

The Master Map File

The master map file (`/etc/auto.master`) is the default configuration file for the `autofs` service. You can use the `/etc/autofs.conf` file to change the master map file for the `autofs` service. Use the `/etc/auto.master.d` directory to configure the master map file. This file identifies the base directory for mount points, and identifies the mapping file to create the automounts.

The name of the master map file is mostly arbitrary (although typically meaningful), and it must have an extension of `.autofs` for the subsystem to recognize it. You can place multiple entries in a single master map file; alternatively, you can create multiple master map files, each with its own logically grouped entries.

The master map uses the following format for automounting:

```
[root@host ~]# cat /etc/auto.master.d/demo.autofs
mount-point    map-file
```

Replace the *mount-point* variable with the directory to be used as the base mount point for automounts.

Replace the *map-file* variable with the file to be used as the map file. The map file must be created before starting the `autofs` service.

The Map File

Map files configure the properties of individual on-demand mount points.

The automounter creates the directories if they do not exist. If the directories already exist before the automounter starts, then the automounter will not remove them when it exits. If a timeout is

specified, then the directory is automatically unmounted if the directory is not accessed for the timeout period.

If you use a single directory name for the mount point, then the directory is mounted as an indirect mount. If you use the full path for the mount point, then the directory is mounted as a direct mount.

The map file uses the following format for automounting:

```
[root@host ~]# cat /etc/auto.demo
mount-point      mount-options    source-location
```

Replace the *mount-point* variable with the `autofs` mount point.

Replace the *mount-options* variable with the options to be used for mounting the `autofs` mount point.

Replace the *source-location* variable with the source location of the mount.

Create an Indirect Map File

Each mapping file identifies the mount point, mount options, and the source location to mount for a set of automounts.

For indirect automounting, include the following content in the master map file:

```
[root@host ~]# cat /etc/auto.master.d/indirect.autofs
/shares /etc/auto.indirect
```

This entry uses the `/shares` directory as the base mount point for the indirect automount. The `/etc/auto.indirect` map file contains the mount details.

The following example shows a sample entry in the map file:

```
[root@host ~]# cat /etc/auto.indirect
work -rw,sync hosta:/shares/work
```

The naming convention for a map file is `/etc/auto.name`, where *name* reflects the content of the map.

Known as the *key* in the man pages, the `autofs` service automatically creates and removes the mount point. In this case, the fully qualified mount point is `/shares/work` (see the master map file). The `autofs` service creates and removes the `/shares` and `/shares/work` directories as needed.

In this example, the local mount point mirrors the server's directory structure. However, this mirroring is not required; the local mount point can have an arbitrary name. The `autofs` service does not enforce a specific naming structure on the client.

Mount options start with a dash character (-) and are comma-separated with no white space. The file-system mount options for manual mounting are also available when automounting. In this example, the automounter mounts the export with read/write access (`rw` option), and the server is synchronized immediately during write operations (`sync` option).

Useful automounter-specific options include `-fstype=` and `-strict`. Use `fstype` to specify the file-system type, for example `nfs4` or `xfs`, and use `strict` to treat errors when mounting file systems as fatal.

The source location for NFS exports follows the `host:/pathname` pattern, in this example `hosta:/shares/work`. For this automount to succeed, the NFS server, `hosta`, must export the directory with read/write access, and the user that requests access must have standard Linux file permissions on the directory. If `hosta` exports the directory with read/only access, then the client gets read/only access even if it requested read/write access.

Create a Direct Map File

A direct map file maps an NFS export to an absolute path mount point. Only one direct map file is necessary, and can contain any number of direct maps.

To use directly mapped mount points, include the following content in the master map file:

```
[root@host ~]# cat /etc/direct.autofs
/- /etc/auto.direct
```

The naming convention for a map file is `/etc/auto.name`, where `name` reflects the content of the map.

All direct map entries use the `/-` path as the base directory. In this case, the `/etc/auto.direct` mapping file contains the mount details.

The following example shows a sample entry in the map file:

```
[root@host ~]# cat /etc/auto.direct
/mnt/docs -rw, sync hosta:/shares/docs
```

The mount point (or key) is always an absolute path. The rest of the mapping file uses the same structure.

In this example, the `/mnt` directory exists, and the `autofs` service does not manage it. The `autofs` service creates and removes the full `/mnt/docs` directory automatically.

Using Wildcards in Map Files

When an NFS server exports multiple subdirectories within a directory, then the automounter can be configured to access any of those subdirectories with a single mapping entry.

Continuing the previous example, if the `hosta:/shares` export contains two or more subdirectories, and they are accessible with the same mount options, then the entry for the `/etc/auto.indirect` file might appear as follows:

```
[root@host ~]# cat /etc/auto.indirect
* -rw, sync hosta:/shares/&
```

The mount point (or key) is an asterisk character (`*`), and the subdirectory on the source location is an ampersand character (`&`). Everything else in the entry is the same.

When a user attempts to access `/shares/work`, the `*` key (which is `work` in this example) replaces the ampersand in the source location and `hosta:/shares/work` is mounted. As

with the indirect map file example, the `autofs` service creates and removes the `work` directory automatically.

Start the Automounter Service

After you have configured the master map and map files, use the `systemctl` command to start and enable the `autofs` service.

```
[root@host ~]# systemctl enable --now autofs
Created symlink /etc/systemd/system/multi-user.target.wants/autofs.service → /usr/lib/systemd/system/autofs.service.
```

The Alternative Automount Method

The `systemd` daemon can automatically create unit files for entries in the `/etc/fstab` file that include the `x-systemd.automount` option. Use the `systemctl daemon-reload` command after modifying an entry's mount options to generate a new unit file. Then use the `systemctl start unit.automount` command to enable the automount configuration.

The naming of the unit is based on its mount location. For example, if the mount point is `/remote/finance`, then the unit file is named `remote-finance.automount`. The `systemd` daemon mounts the file system when the `/remote/finance` directory is initially accessed.

This method can be simpler than installing and configuring the `autofs` service. However, a `systemd.automount` unit can support only absolute path mount points, similar to the `autofs` direct maps.



References

`autofs(5)`, `automount(8)`, `auto.master(5)`, `mount.nfs(8)`, and `systemd.automount(5)` man pages

For further information, refer to the *Mounting File Systems on Demand* chapter in the *Managing File Systems* guide at

https://docs.redhat.com/en/documentation/red_hat_enterprise_linux/9/html-single/managing_file_systems/index#the-autofs-service_mounting-file-systems-on-demand

► Guided Exercise

Automount Network-Attached Storage

Create direct-mapped and indirect-mapped automount-managed mount points that mount NFS file systems.

Outcomes

- Install required packages for the automounter.
- Configure direct and indirect automounter maps, with resources from a preconfigured NFSv4 server.
- Describe the difference between direct and indirect automounter maps.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This start script determines whether `servera` and `serverb` are reachable on the network. The script alerts you if those servers are not available. The start script configures `serverb` as an NFSv4 server, sets up permissions, and exports directories. The script also creates users and groups that are needed on both `servera` and `serverb`.

```
[student@workstation ~]$ lab start netstorage-autofs
```

Instructions

An internet service provider uses a central server, `serverb`, to host shared directories with important documents that must be available on demand. When users log in to `servera`, they need access to the automounted shared directories.

The following list provides the environment characteristics for completing this exercise:

- The `serverb` machine exports the `/shares/indirect` directory, which in turn contains the `west`, `central`, and `east` subdirectories.
- The `serverb` machine also exports the `/shares/direct/external` directory.
- The `operators` group consists of the `operator1` and `operator2` users. They have read and write access to the `/shares/indirect/west`, `/shares/indirect/central`, and `/shares/indirect/east` exported directories.
- The `contractors` group consists of the `contractor1` and `contractor2` users. They have read and write access to the `/shares/direct/external` exported directory.
- The expected mount points for `servera` are `/external` and `/internal`.
- The `/shares/direct/external` exported directory is automounted on `servera` with a `direct` map on `/external`.

- The `/shares/indirect/west` exported directory is automounted on `servera` with an `indirect` map on `/internal/west`.
- The `/shares/indirect/central` exported directory is automounted on `servera` with an `indirect` map on `/internal/central`.
- The `/shares/indirect/east` exported directory is automounted on `servera` with an `indirect` map on `/internal/east`.
- All user passwords are set to `redhat`.
- The `nfs-utils` package is already installed.

► 1. Log in to `servera` and install the required packages.

- 1.1. Log in to `servera` as the `student` user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

- 1.2. Install the `autofs` package.

```
[root@servera ~]# dnf install autofs
...output omitted...
Is this ok [y/N]: y
...output omitted...
Complete!
```

► 2. Configure an automounter direct map on `servera` with exports from `serverb`. Create the direct map with files that are named `/etc/auto.master.d/direct.autofs` for the master map and `/etc/auto.direct` for the mapping file. Use the `/external` directory as the main mount point on `servera`.

- 2.1. Test the NFS server and export before you configure the automounter.

```
[root@servera ~]# mount -t nfs \
serverb.lab.example.com:/shares/direct/external /mnt
[root@servera ~]# ls -l /mnt
total 4
-rw-r--r--. 1 root contractors 22 Apr  7 23:15 README.txt
[root@servera ~]# umount /mnt
```

- 2.2. Create a master map file named `/etc/auto.master.d/direct.autofs`, insert the following content, and save the changes.

```
/- /etc/auto.direct
```

- 2.3. Create a direct map file named `/etc/auto.direct`, insert the following content, and save the changes.

```
/external -rw, sync, fstype=nfs4 serverb.lab.example.com:/shares/direct/external
```

- ▶ **3.** Configure an automounter indirect map on `servera` with exports from `serverb`. Create the indirect map with files that are named `/etc/auto.master.d/indirect.autofs` for the master map and `/etc/auto.indirect` for the mapping file. Use the `/internal` directory as the main mount point on `servera`.

3.1. Test the NFS server and export before you configure the automounter.

```
[root@servera ~]# mount -t nfs \
serverb.lab.example.com:/shares/indirect /mnt
[root@servera ~]# ls -l /mnt
total 0
drwxrws---. 2 root operators 24 Apr  7 23:34 central
drwxrws---. 2 root operators 24 Apr  7 23:34 east
drwxrws---. 2 root operators 24 Apr  7 23:34 west
[root@servera ~]# umount /mnt
```

3.2. Create a master map file named `/etc/auto.master.d/indirect.autofs`, insert the following content, and save the changes.

```
/internal /etc/auto.indirect
```

3.3. Create an indirect map file named `/etc/auto.indirect`, insert the following content, and save the changes.

```
* -rw, sync, fstype=nfs4 serverb.lab.example.com:/shares/indirect/&
```

- ▶ **4.** Start the `autofs` service on `servera`, and enable it to start automatically at boot time.

4.1. Start and enable the `autofs` service on `servera`.

```
[root@servera ~]# systemctl enable --now autofs
Created symlink /etc/systemd/system/multi-user.target.wants/autofs.service → /usr/lib/systemd/system/autofs.service.
```

- ▶ **5.** Test the direct automounter map as the `contractor1` user. When done, exit from the `contractor1` user session on `servera`.

5.1. Switch to the `contractor1` user.

```
[root@servera ~]# su - contractor1
[contractor1@servera ~]$
```

5.2. List the `/external` mount point.

```
[contractor1@servera ~]$ ls -l /external
total 4
-rw-r--r--. 1 root contractors 22 Apr  7 23:34 README.txt
```

5.3. Review the content and test the access to the `/external` mount point.

```
[contractor1@servera ~]$ cat /external/README.txt
###External Folder###
[contractor1@servera ~]$ echo testing-direct > /external/testing.txt
[contractor1@servera ~]$ cat /external/testing.txt
testing-direct
```

5.4. Exit from the `contractor1` user session.

```
[contractor1@servera ~]$ exit
logout
[root@servera ~]#
```

- 6. Test the indirect automounter map as the `operator1` user. When done, log out from `servera`.

6.1. Switch to the `operator1` user.

```
[root@servera ~]# su - operator1
[operator1@servera ~]$
```

6.2. List the `/internal` mount point.

```
[operator1@servera ~]$ ls -l /internal
total 0
```



Note

With an automounter indirect map, you must access each exported subdirectory for them to mount. With an automounter direct map, after you access the mapped mount point, you can immediately view and access the subdirectories and content in the exported directory.

6.3. Test the `/internal/west` automounter exported directory access.

```
[operator1@servera ~]$ ls -l /internal/west/
total 4
-rw-r--r--. 1 root operators 18 Apr  7 23:34 README.txt
[operator1@servera ~]$ cat /internal/west/README.txt
###West Folder###
[operator1@servera ~]$ echo testing-1 > /internal/west/testing-1.txt
[operator1@servera ~]$ cat /internal/west/testing-1.txt
testing-1
[operator1@servera ~]$ ls -l /internal
total 0
drwxrws---. 2 root operators 24 Apr  7 23:34 west
```

6.4. Test the `/internal/central` automounter exported directory access.

```
[operator1@servera ~]$ ls -l /internal/central
total 4
-rw-r--r--. 1 root operators 21 Apr  7 23:34 README.txt
[operator1@servera ~]$ cat /internal/central/README.txt
###Central Folder###
[operator1@servera ~]$ echo testing-2 > /internal/central/testing-2.txt
[operator1@servera ~]$ cat /internal/central/testing-2.txt
testing-2
[operator1@servera ~]$ ls -l /internal
total 0
drwxrws---. 2 root operators 24 Apr  7 23:34 central
drwxrws---. 2 root operators 24 Apr  7 23:34 west
```

6.5. Test the /internal/east automounter exported directory access.

```
[operator1@servera ~]$ ls -l /internal/east
total 4
-rw-r--r--. 1 root operators 18 Apr  7 23:34 README.txt
[operator1@servera ~]$ cat /internal/east/README.txt
###East Folder###
[operator1@servera ~]$ echo testing-3 > /internal/east/testing-3.txt
[operator1@servera ~]$ cat /internal/east/testing-3.txt
testing-3
[operator1@servera ~]$ ls -l /internal
total 0
drwxrws---. 2 root operators 24 Apr  7 23:34 central
drwxrws---. 2 root operators 24 Apr  7 23:34 east
drwxrws---. 2 root operators 24 Apr  7 23:34 west
```

6.6. Test the /external automounter exported directory access.

```
[operator1@servera ~]$ ls -l /external
ls: cannot open directory '/external': Permission denied
```

6.7. Return to the workstation machine as the student user.

```
[operator1@servera ~]$ exit
logout
[root@servera ~]# exit
logout
[student@servera ~]$ exit
logout
Connection to servera closed.
```

Finish

On the workstation machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish netstorage-autofs
```


► Lab

Access Network-Attached Storage

Configure the automounter with an indirect map, using exports from an NFSv4 server.

Outcomes

- Install required packages to set up the automounter.
- Configure an automounter indirect map, with resources from a preconfigured NFSv4 server.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This start script determines whether the `servera` and `serverb` systems are reachable on the network. The start script configures `serverb` as an NFSv4 server, sets up permissions, and exports directories. The script also creates users and groups that are needed on both `servera` and `serverb` systems.

```
[student@workstation ~]$ lab start netstorage-review
```

Instructions

An IT support company uses a central server, `serverb`, to host some exported directories on `/shares` for their groups and users. Users must be able to log in and have their exported directories mounted on demand and ready to use, in the `/remote` directory on `servera`.

The following list provides the environment characteristics for completing this exercise:

- The `serverb` machine is sharing the `/shares` directory, which in turn contains the `management`, `production`, and `operation` subdirectories.
- The `manager`s group consists of the `manager1` and `manager2` users. Those users have read and write access to the `/shares/management` exported directory.
- The `production` group consists of the `dbuser1` and `sysadmin1` users. Those users have read and write access to the `/shares/production` exported directory.
- The `operator`s group consists of the `contractor1` and `consultant1` users. Those users have read and write access to the `/shares/operation` exported directory.
- The main mount point for `servera` is the `/remote` directory.
- Use the `/etc/auto.master.d/shares.autofs` file as the master map file, and use the `/etc/auto.shares` file as the indirect map file.
- The `/shares/management` exported directory is automounted on `/remote/management` on `servera`.

- The `/shares/production` exported directory is automounted on `/remote/production` on `servera`.
 - The `/shares/operation` exported directory is automounted on `/remote/operation` on `servera`.
 - All user passwords are set to `redhat`.
1. Log in to `servera` and install the required packages.
 2. Configure an automounter indirect map on `servera` with exports from `serverb`. Create an indirect map with files that are named `/etc/auto.master.d/shares.autofs` for the master map and `/etc/auto.shares` for the mapping file. Use the `/remote` directory as the main mount point on `servera`. Reboot `servera` to determine whether the `autofs` service starts automatically.
 3. Test the `autofs` configuration with the various users. When done, log out from `servera`.

Evaluation

On the `workstation` machine, use the `lab` command to confirm success of this exercise.

```
[student@workstation ~]$ lab grade netstorage-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish netstorage-review
```

► Solution

Access Network-Attached Storage

Configure the automounter with an indirect map, using exports from an NFSv4 server.

Outcomes

- Install required packages to set up the automounter.
- Configure an automounter indirect map, with resources from a preconfigured NFSv4 server.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This start script determines whether the `servera` and `serverb` systems are reachable on the network. The start script configures `serverb` as an NFSv4 server, sets up permissions, and exports directories. The script also creates users and groups that are needed on both `servera` and `serverb` systems.

```
[student@workstation ~]$ lab start netstorage-review
```

Instructions

An IT support company uses a central server, `serverb`, to host some exported directories on `/shares` for their groups and users. Users must be able to log in and have their exported directories mounted on demand and ready to use, in the `/remote` directory on `servera`.

The following list provides the environment characteristics for completing this exercise:

- The `serverb` machine is sharing the `/shares` directory, which in turn contains the `management`, `production`, and `operation` subdirectories.
- The `manager`s group consists of the `manager1` and `manager2` users. Those users have read and write access to the `/shares/management` exported directory.
- The `production` group consists of the `dbuser1` and `sysadmin1` users. Those users have read and write access to the `/shares/production` exported directory.
- The `operator`s group consists of the `contractor1` and `consultant1` users. Those users have read and write access to the `/shares/operation` exported directory.
- The main mount point for `servera` is the `/remote` directory.
- Use the `/etc/auto.master.d/shares.autofs` file as the master map file, and use the `/etc/auto.shares` file as the indirect map file.
- The `/shares/management` exported directory is automounted on `/remote/management` on `servera`.

- The `/shares/production` exported directory is automounted on `/remote/production` on servera.
- The `/shares/operation` exported directory is automounted on `/remote/operation` on servera.
- All user passwords are set to `redhat`.

1. Log in to servera and install the required packages.

- 1.1. Log in to servera as the student user and switch to the root user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

- 1.2. Install the `autofs` package.

```
[root@servera ~]# dnf install autofs
...output omitted...
Is this ok [y/N]: y
...output omitted...
Complete!
```

2. Configure an automounter indirect map on servera with exports from serverb. Create an indirect map with files that are named `/etc/auto.master.d/shares.autofs` for the master map and `/etc/auto.shares` for the mapping file. Use the `/remote` directory as the main mount point on servera. Reboot servera to determine whether the `autofs` service starts automatically.

- 2.1. Test the NFS server before you configure the automounter.

```
[root@servera ~]# mount -t nfs serverb.lab.example.com:/shares /mnt
[root@servera ~]# ls -l /mnt
total 0
drwxrwx---. 2 root managers 25 Apr 4 01:13 management
drwxrwx---. 2 root operators 25 Apr 4 01:13 operation
drwxrwx---. 2 root production 25 Apr 4 01:13 production
[root@servera ~]# umount /mnt
```

- 2.2. Create a master map file named `/etc/auto.master.d/shares.autofs`, insert the following content, and save the changes.

```
/remote /etc/auto.shares
```

- 2.3. Create an indirect map file named `/etc/auto.shares`, insert the following content, and save the changes.

```
* -rw, sync, fstype=nfs4 serverb.lab.example.com:/shares/&
```

- 2.4. Start and enable the `autofs` service on servera.

```
[root@servera ~]# systemctl enable --now autofs
Created symlink /etc/systemd/system/multi-user.target.wants/autofs.service → /usr/lib/systemd/system/autofs.service.
```

3. Test the autofs configuration with the various users. When done, log out from servera.

3.1. Switch to the manager1 user and test access.

```
[root@servera ~]# su - manager1
[manager1@servera ~]$ ls -l /remote/management/
total 4
-rw-r--r--. 1 root managers 46 Apr  4 01:13 Welcome.txt
[manager1@servera ~]$ cat /remote/management/Welcome.txt
###Welcome to Management Folder on SERVERB###
[manager1@servera ~]$ echo TEST1 > /remote/management/Test.txt
[manager1@servera ~]$ cat /remote/management/Test.txt
TEST1
[manager1@servera ~]$ ls -l /remote/operation/
ls: cannot open directory '/remote/operation/': Permission denied
[manager1@servera ~]$ ls -l /remote/production/
ls: cannot open directory '/remote/production/': Permission denied
[manager1@servera ~]$ exit
logout
[root@servera ~]#
```

3.2. Switch to the dbuser1 user and test access.

```
[root@servera ~]# su - dbuser1
[dbuser1@servera ~]$ ls -l /remote/production/
total 4
-rw-r--r--. 1 root production 46 Apr  4 01:13 Welcome.txt
[dbuser1@servera ~]$ cat /remote/production/Welcome.txt
###Welcome to Production Folder on SERVERB###
[dbuser1@servera ~]$ echo TEST2 > /remote/production/Test.txt
[dbuser1@servera ~]$ cat /remote/production/Test.txt
TEST2
[dbuser1@servera ~]$ ls -l /remote/operation/
ls: cannot open directory '/remote/operation/': Permission denied
[dbuser1@servera ~]$ ls -l /remote/management/
ls: cannot open directory '/remote/management/': Permission denied
[dbuser1@servera ~]$ exit
logout
[root@servera ~]#
```

3.3. Switch to the contractor1 user and test access.

```
[root@servera ~]# su - contractor1
[contractor1@servera ~]$ ls -l /remote/operation/
total 4
-rw-r--r--. 1 root operators 45 Apr  4 01:13 Welcome.txt
[contractor1@servera ~]$ cat /remote/operation/Welcome.txt
###Welcome to Operation Folder on SERVERB###
```

```
[contractor1@servera ~]$ echo TEST3 > /remote/operation/Test.txt
[contractor1@servera ~]$ cat /remote/operation/Test.txt
TEST3
[contractor1@servera ~]$ ls -l /remote/management/
ls: cannot open directory '/remote/management/': Permission denied
[contractor1@servera ~]$ ls -l /remote/production/
ls: cannot open directory '/remote/production/': Permission denied
[contractor1@servera ~]$ exit
logout
[root@servera ~]#
```

3.4. Explore the mount options for the NFS automounted export.

```
[root@servera ~]# mount | grep nfs
rpc_pipefs on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw,relatime)
serverb.lab.example.com:/shares/management on /remote/management type nfs4
(rw,relatime,vers=4.2,rsz=262144,wsz=262144,namlen=255,
sync,proto=tcp,timeo=600,retrans=2,sec=sys,clientaddr=172.25.250.10,
local_lock=none,addr=172.25.250.11)
serverb.lab.example.com:/shares/operation on /remote/operation type nfs4
(rw,relatime,vers=4.2,rsz=262144,wsz=262144,namlen=255,
sync,proto=tcp,timeo=600,retrans=2,sec=sys,clientaddr=172.25.250.10,
local_lock=none,addr=172.25.250.11)
serverb.lab.example.com:/shares/production on /remote/production type nfs4
(rw,relatime,vers=4.2,rsz=262144,wsz=262144,namlen=255,
sync,proto=tcp,timeo=600,retrans=2,sec=sys,clientaddr=172.25.250.10,
local_lock=none,addr=172.25.250.11)
```

3.5. Return to the workstation machine as the student user.

```
[root@servera ~]# exit
logout
[student@servera ~]$ exit
logout
Connection to servera closed.
```

Evaluation

On the workstation machine, use the `lab` command to confirm success of this exercise.

```
[student@workstation ~]$ lab grade netstorage-review
```

Finish

On the workstation machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish netstorage-review
```

Summary

- Mount and unmount an NFS share from the command line.
- Configure an NFS share to mount automatically at startup.
- Configure the automounter with direct and indirect maps, and describe their differences.

Chapter 10

Control the Boot Process

Goal

Manage the boot process to control offered services and to troubleshoot and repair problems.

Objectives

- Describe the Red Hat Enterprise Linux boot process, set the default target when booting, and boot a system to a non-default target.
- Log in to a system and change the root password when the current root password is lost.
- Manually repair file-system configuration or corruption issues that stop the boot process.

Sections

- Select the Boot Target (and Guided Exercise)
- Reset the Root Password (and Guided Exercise)
- Repair File-system Issues at Boot (and Guided Exercise)

Lab

- Control the Boot Process

Select the Boot Target

Objectives

- Describe the Red Hat Enterprise Linux boot process, set the default target when booting, and boot a system to a non-default target.

Describe the Red Hat Enterprise Linux 9 Boot Process

Modern computer systems are complex combinations of hardware and software. Starting from an undefined, powered-down state to a running system with a login prompt requires many pieces of hardware and software to work together. The following list gives a high-level overview of the tasks for a physical x86_64 system that boots Red Hat Enterprise Linux 9. The list for x86_64 virtual machines is similar, except that the hypervisor handles some hardware-specific steps in software.

- The machine is powered on. The system firmware, either modern UEFI or earlier BIOS, runs a *Power On Self Test (POST)* and starts to initialize the hardware.

The system BIOS or UEFI is configured by pressing a specific key combination, such as F2, early during the boot process.

- The UEFI boot firmware is configured by searching for a bootable device, which searches for or configures the *Master Boot Record (MBR)* on all disks.

The system BIOS or UEFI configuration is configured by pressing a specific key combination, such as F2, early during the boot process.

- The system firmware reads a boot loader from disk and then passes control of the system to the boot loader. On a Red Hat Enterprise Linux 9 system, the boot loader is the *GRand Unified Bootloader version 2 (GRUB2)*.

The `grub2-install` command installs GRUB2 as the boot loader on the disk for BIOS systems. Do not use the `grub2-install` command directly to install the UEFI boot loader. RHEL 9 provides a prebuilt `/boot/efi/EFI/redhat/grubx64.efi` file, which contains the required authentication signatures for a Secure Boot system. Executing `grub2-install` directly on a UEFI system generates a new `grubx64.efi` file without those required signatures. You can restore the correct `grubx64.efi` file from the `grub2-efi` package.

- GRUB2 loads its configuration from the `/boot/grub2/grub.cfg` file for BIOS, and from the `/boot/efi/EFI/redhat/grub.cfg` file for UEFI, and displays a menu to select which kernel to boot.

GRUB2 is configured by using the `/etc/grub.d/` directory and the `/etc/default/grub` file. The `grub2-mkconfig` command generates the `/boot/grub2/grub.cfg` or `/boot/efi/EFI/redhat/grub.cfg` files for BIOS or UEFI, respectively.

- After you select a kernel, or the timeout expires, the boot loader loads the kernel and *initramfs* from disk and places them in memory. An *initramfs* image is an archive with the kernel modules for all the required hardware at boot, initialization scripts, and more. In Red Hat Enterprise Linux 9, the *initramfs* image contains a bootable root file system with a running kernel and a `systemd` unit.

The `initramfs` image is configured by using the `/etc/dracut.conf.d/` directory, the `dracut` command, and the `lsinitrd` command to inspect the `initramfs` file.

- The boot loader hands control over to the kernel, and passes in any specified options on the kernel command line in the boot loader, and the location of the `initramfs` image in memory.

The boot loader is configured by using the `/etc/grub.d/` directory, the `/etc/default/grub` file, and the `grub2-mkconfig` command to generate the `/boot/grub2/grub.cfg` file.

- The kernel initializes all hardware for which it can find a driver in the `initramfs` image, and then executes the `/sbin/init` script from the `initramfs` image as PID 1. On Red Hat Enterprise Linux 9, the `/sbin/init` script is a link to the `systemd` unit.

The script is configured by using the kernel `init=` command-line parameter.

- The `systemd` unit from the `initramfs` image executes all units for the `initrd.target` target. This unit includes mounting the root file system on disk to the `/sysroot` directory.

Configured by using the `/etc/fstab` file.

- The kernel switches (pivots) the root file system from the `initramfs` image to the root file system in the `/sysroot` directory. The `systemd` unit then re-executes itself by using the installed copy of the `systemd` unit on the disk.
- The `systemd` unit looks for a default target, which is either passed in from the kernel command line or is configured on the system. The `systemd` unit then starts (and stops) units to comply with the configuration for that target, and solves dependencies between units automatically. A `systemd` target is a set of units that the system activates to reach the intended state. These targets typically start a text-based login or a graphical login screen.

Configured by using the `/etc/systemd/system/default.target` file and the `/etc/systemd/system/` directory.




Figure 10.1: Boot process for BIOS-based and UEFI-based systems

Power Off and Reboot

To power off or reboot a running system from the command line, you can use the `systemctl` command.

The `systemctl poweroff` command stops all running services, unmounts all file systems (or remounts them read-only when they cannot be unmounted), and then powers down the system.

The `systemctl reboot` command stops all running services, unmounts all file systems, and then reboots the system.

You can also use the shorter version of these commands, `poweroff` and `reboot`, which are symbolic links to their `systemctl` equivalents.

Note

The `systemctl halt` and `halt` commands are also available to stop the system. Unlike the `poweroff` command, these commands do not power off the system; they bring down a system to a point where it is safe to power it off manually.

Select a Systemd Target

A `systemd` target is a set of `systemd` units that the system must start to reach an intended state. The following table lists the most important targets:

Commonly Used Targets

Target	Purpose
<code>graphical.target</code>	This target supports multiple users, and provides graphical- and text-based logins.
<code>multi-user.target</code>	This target supports multiple users, and provides text-based logins only.
<code>rescue.target</code>	This target provides a single-user system to enable repairing your system.
<code>emergency.target</code>	This target starts the most minimal system for repairing your system when the <code>rescue.target</code> unit fails to start.

A target can be a part of another target. For example, the `graphical.target` unit includes the `multi-user.target` unit, which in turn depends on the `basic.target` unit and others. You can view these dependencies with the following command:

```
[user@host ~]$ systemctl list-dependencies graphical.target | grep target
graphical.target
* └─multi-user.target
*   └─basic.target
*     └─paths.target
*       └─slices.target
*         └─sockets.target
*           └─sysinit.target
*             └─cryptsetup.target
```

```
* | | | integritysetup.target
* | | | local-fs.target
...output omitted...
```

To list the available targets, use the following command:

```
[user@host ~]$ systemctl list-units --type=target --all
UNIT                                LOAD    ACTIVE SUB    DESCRIPTION
-----
basic.target                       loaded active active Basic System
...output omitted...
cloud-config.target               loaded active active Cloud-config availability
cloud-init.target                 loaded active active Cloud-init target
cryptsetup-pre.target             loaded inactive dead Local Encrypted Volumes
(Pre)
cryptsetup.target                 loaded active active Local Encrypted Volumes
...output omitted...
```

Select a Target at Runtime

On a running system, administrators can switch to a different target by using the `systemctl isolate` command.

```
[root@host ~]# systemctl isolate multi-user.target
```

Isolating a target stops all services that the target does not require (and its dependencies), and starts any required services that are not yet started.

Not all targets can be isolated. You can isolate only targets where `AllowIsolate=yes` is set in their unit files. For example, you can isolate the graphical target, but not the `cryptsetup` target.

```
[user@host ~]$ systemctl cat graphical.target
# /usr/lib/systemd/system/graphical.target
...output omitted...
[Unit]
Description=Graphical Interface
Documentation=man:systemd.special(7)
Requires=multi-user.target
Wants=display-manager.service
Conflicts=rescue.service rescue.target
After=multi-user.target rescue.service rescue.target display-manager.service
AllowIsolate=yes
[user@host ~]$ systemctl cat cryptsetup.target
# /usr/lib/systemd/system/cryptsetup.target
...output omitted...
[Unit]
Description=Local Encrypted Volumes
Documentation=man:systemd.special(7)
```

Set a Default Target

When the system starts, the `systemd` unit activates the `default.target` target. Normally, the default target the `/etc/systemd/system/` directory is a symbolic link to either the

`graphical.target` or the `multi-user.target` targets. Instead of editing this symbolic link by hand, the `systemctl` command provides two subcommands to manage this link: `get-default` and `set-default`.

```
[root@host ~]# systemctl get-default
multi-user.target
[root@host ~]# systemctl set-default graphical.target
Removed /etc/systemd/system/default.target.
Created symlink /etc/systemd/system/default.target -> /usr/lib/systemd/system/
graphical.target.
[root@host ~]# systemctl get-default
graphical.target
```

Select a Different Target at Boot Time

To select a different target at boot time, append the `systemd.unit=target.target` option to the kernel command line from the boot loader.

For example, to boot the system into a rescue shell where you can change the system configuration with almost no services running, append the following option to the kernel command line from the boot loader:

```
systemd.unit=rescue.target
```

This configuration change affects only a single boot, and is a useful tool to troubleshoot the boot process.

To use this method to select a different target, use the following procedure:

1. Boot or reboot the system.
2. Interrupt the boot loader menu countdown by pressing any key (except `Enter`, which would initiate a normal boot).
3. Move the cursor to the kernel entry to start.
4. Press `e` to edit the current entry.
5. Move the cursor to the line that starts with `linux` which is the kernel command line.
6. Append `systemd.unit=target.target`, for example, `systemd.unit=emergency.target`.
7. Press `Ctrl+x` to boot with these changes.



References

`info grub2` (*GNU GRUB manual*)

`bootup(7)`, `dracut.bootup(7)`, `lsinitrd(1)`, `systemd.target(5)`, `systemd.special(7)`, `sulogin(8)`, and `systemctl(1)` man pages

For more information, refer to the *Managing Services with systemd* chapter in the *Configuring Basic System Settings* guide at https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/configuring_basic_system_settings/index#managing-services-with-systemd

► Guided Exercise

Select the Boot Target

Determine the default target into which a system boots, and boot that system into other targets.

Outcomes

- Update the system default target and use a temporary target from the boot loader.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command ensures that all required resources are available.

```
[student@workstation ~]$ lab start boot-selecting
```

Instructions

- 1. On the `workstation` machine, open a terminal and confirm that the default target is `graphical.target`.

```
[student@workstation ~]$ systemctl get-default
graphical.target
```

- 2. On the `workstation` machine, switch to the `multi-user` target manually without rebooting. Use the `sudo` command and if prompted, use `student` as the password.

```
[student@workstation ~]$ sudo systemctl isolate multi-user.target
[sudo] password for student: student
```

- 3. Log in as the `root` user by using `redhat` as the password.

```
workstation login: root
Password: redhat
[root@workstation ~]#
```

- 4. Configure the `workstation` machine to automatically boot into the `multi-user` target, and then reboot the `workstation` machine to verify. When done, change the default `systemd` target back to the `graphical` target.

- 4.1. Set the default target.

```
[root@workstation ~]# systemctl set-default multi-user.target
Removed /etc/systemd/system/default.target.
Created symlink /etc/systemd/system/default.target -> /usr/lib/systemd/system/
multi-user.target.
```

- 4.2. Reboot the `workstation` machine. After reboot, the system presents a text-based console and not a graphical login screen.

```
[root@workstation ~]# systemctl reboot
```

- 4.3. Log in as the `root` user.

```
workstation login: root
Password: redhat
Last login: Thu Mar 28 14:50:53 on tty1
[root@workstation ~]#
```

- 4.4. Set the default `systemd` target back to the graphical target.

```
[root@workstation ~]# systemctl set-default graphical.target
Removed /etc/systemd/system/default.target.
Created symlink /etc/systemd/system/default.target -> /usr/lib/systemd/system/
graphical.target.
```

This step concludes the first part of the exercise, where you practice setting the default `systemd` target.

- ▶ 5. In this second part of the exercise, you practice by using rescue mode to recover the system.
Access the boot loader by rebooting `workstation` again. From within the boot loader menu, boot into the rescue target.

- 5.1. Initiate the reboot.

```
[root@workstation ~]# systemctl reboot
```

- 5.2. When the boot loader menu appears, press any key to interrupt the countdown (except `Enter`, which would initiate a normal boot).
- 5.3. Use the cursor keys to highlight the default boot loader entry.
- 5.4. Press `e` to edit the current entry.
- 5.5. Using the cursor keys, navigate to the line that starts with `linux`.
- 5.6. Press `End` to move the cursor to the end of the line.
- 5.7. Append `systemd.unit=rescue.target` to the end of the line.

**Note**

If it is difficult for you to read the text in the console, then consider changing the resolution when you edit the kernel line in the boot loader entry.

To change the console resolution, add either `video=640x480` or `vga=ask` to the line that starts with the `linux` word, after `systemd.unit=rescue.target`. With the `video=640x480` argument, the text console should be displayed at approximately an 80-column width. If you use `vga=ask` instead, then you can choose the resolution that best suits your environment at boot time.

5.8. Press `Ctrl+x` to boot by using the modified configuration.

5.9. Log in to rescue mode. You might need to press `Enter` to get a clean prompt.

```
Give root password for maintenance
(or press Control-D to continue): redhat
[root@workstation ~]#
```

- ▶ 6. Confirm that in rescue mode, the root file system is in read/write mode.

```
[root@workstation ~]# mount
...output omitted...
/dev/vda4 on / type xfs
(rw,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,noquota)
...output omitted...
```

- ▶ 7. Press `Ctrl+d` to continue with the boot process.

The system presents a graphical login. Log in as the `student` user.

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish boot-selecting
```

Reset the Root Password

Objectives

- Log in to a system and change the root password when the current root password is lost.

Reset the Root Password from the Boot Loader

One task that every system administrator should be able to accomplish is resetting a lost `root` password. This task is trivial if the administrator is still logged in, either as an unprivileged user but with full `sudo` access, or as `root`. This task becomes slightly more involved when the administrator is not logged in.

Several methods exist to set a new `root` password. A system administrator could, for example, boot the system by using a Live CD, mount the root file system from there, and edit `/etc/shadow`. This section explores a method that does not require the use of external media.

On Red Hat Enterprise Linux 9, the scripts that run from the `initramfs` image can be paused at certain points, to provide a `root` shell, and then continue when that shell exits. This script is mostly meant for debugging, and also to reset a lost `root` password.

Starting from Red Hat Enterprise Linux 9, if you install your system from a DVD, then the default kernel asks for the `root` password when you try to enter maintenance mode. Thus, to reset a lost `root` password, you must use the rescue kernel.

To access that `root` shell, follow these steps:

1. Reboot the system.
2. Interrupt the boot-loader countdown by pressing any key, except `Enter`.
3. Move the cursor to the rescue kernel entry to boot (the entry with the *rescue* word in its name).
4. Press `e` to edit the selected entry.
5. Move the cursor to the kernel command line (the line that starts with `linux`).
6. Append `rd.break`. With that option, the system breaks just before the system hands control from the `initramfs` image to the actual system.
7. Press `Ctrl+x` to boot with the changes.
8. Press `Enter` to perform maintenance when prompted.

At this point, the system presents a `root` shell, and the root file system on the disk is mounted read-only on `/sysroot`. Because troubleshooting often requires modifying the root file system, you must remount the root file system as read/write. The following step shows how the `remount, rw` option to the `mount` command remounts the file system where the new option (`rw`) is set.

**Important**

Because the system has not yet enabled SELinux, any file that you create does not have SELinux context. Some tools, such as the `passwd` command, first create a temporary file, and then replace it with the file that is intended for editing, which effectively creates a file without SELinux context. For this reason, when you use the `passwd` command with `rd.break`, the `/etc/shadow` file does not receive SELinux context.

To reset the root password, use the following procedure:

1. Remount `/sysroot` as read/write.

```
sh-5.1# mount -o remount,rw /sysroot
```

2. Switch into a `chroot` jail, where `/sysroot` is treated as the root of the file-system tree.

```
sh-5.1# chroot /sysroot
```

3. Set a new root password.

```
sh-5.1# passwd root
```

4. Ensure that all unlabeled files, including `/etc/shadow` at this point, get relabeled during boot.

```
sh-5.1# touch /.autorelabel
```

5. Type `exit` twice. The first command exits the `chroot` jail, and the second command exits the `initramfs` debug shell.

At this point, the system continues booting, performs a full SELinux relabeling, and then reboots again.

Recovery of a Cloud Image-based System

If your system was installed by deploying and modifying one of the official cloud images instead of using the installer, then some system configuration aspects might differ for the boot process.

The procedure to use the `rd.break` option to get a root shell is similar to the previously outlined procedure, with some minor changes.

If your system was deployed from a Red Hat Enterprise Linux cloud image, then your boot menu does not have a rescue kernel by default. However, you can use the default kernel to enter maintenance mode by using the `rd.break` option without entering the root password.

The kernel prints boot messages and displays the root prompt on the system console. Prebuilt images might have multiple `console=` arguments on the kernel command line in the bootloader. Even though the system sends the kernel messages to all the consoles, the root shell that the `rd.break` option sets up uses the last console that is specified on the command line. If you do not get your prompt, then you might temporarily reorder the `console=` arguments when you edit the kernel command line in the boot loader.

Inspect Logs

Looking at the logs of previously failed boots can be useful. If the system journals persist across reboots, then you can use the `journalctl` tool to inspect those logs.

Remember that by default, the system journals are kept in the `/run/log/journal` directory, and the journals are cleared when the system reboots. To store journals in the `/var/log/journal` directory, which persists across reboots, set the `Storage` parameter to `persistent` in the `/etc/systemd/journald.conf` file.

```
[root@host ~]# vim /etc/systemd/journald.conf
...output omitted...
[Journal]
Storage=persistent
...output omitted...
[root@host ~]# systemctl restart systemd-journald.service
```

To inspect the logs of a previous boot, use the `journalctl` command `-b` option. Without any arguments, the `journalctl` command `-b` option displays only messages since the last boot. With a negative number as an argument, it displays the logs of previous boots.

```
[root@host ~]# journalctl -b -1 -p err
```

This command shows all messages that are rated as an error or worse from the previous boot.

Repair Systemd Boot Issues

To troubleshoot service startup issues at boot time, Red Hat Enterprise Linux 8 and later versions provide the following tools:

Enable the Early Debug Shell

By enabling the `debug-shell.service` with the `systemctl enable debug-shell.service` command, the system spawns a `root` shell on `TTY9` (`Ctrl+Alt+F9`) early during the boot sequence. This shell is automatically logged in as `root`, so that administrators can debug the system when the operating system is still booting.



Warning

Disable the `debug-shell.service` service when you are done debugging, because it leaves an unauthenticated `root` shell open to anyone with local console access.

Alternatively, to activate the debug shell during the boot by using the GRUB2 menu, follow these steps:

1. Reboot the system.
2. Interrupt the boot-loader countdown by pressing any key, except `Enter`.
3. Move the cursor to the kernel entry to boot.
4. Press `e` to edit the selected entry.
5. Move the cursor to the kernel command line (the line that starts with `linux`).

6. Append `systemd.debug-shell`. With this parameter, the system boots into the debug shell.
7. Press `Ctrl+x` to boot with the changes.

Use the Emergency and Rescue Targets

By appending either `systemd.unit=rescue.target` or `systemd.unit=emergency.target` to the kernel command line from the boot loader, the system enters into a rescue or emergency shell instead of starting normally. Both of these shells require the `root` password.

The emergency target keeps the root file system mounted read-only, while the rescue target waits for the `sysinit.target` unit to complete, so that more of the system is initialized, such as the logging service or the file systems. The root user at this point cannot change `/etc/fstab` until the drive is remounted in a read write state with the `mount -o remount,rw /` command.

Administrators can use these shells to fix any issues that prevent the system from booting normally, for example, a dependency loop between services, or an incorrect entry in `/etc/fstab`. Exiting from these shells continues with the regular boot process.

Identify Stuck Jobs

During startup, `systemd` spawns various jobs. If some of these jobs cannot complete, then they block other jobs from running. To inspect the current job list, administrators can use the `systemctl list-jobs` command. Any jobs that are listed as running must complete before the jobs that are listed as waiting can continue.



References

`dracut.cmdline(7)`, `systemd-journald(8)`, `journald.conf(5)`, `journaldctl(1)`, and `systemctl(1)` man pages

► Guided Exercise

Reset the Root Password

Reset the root password on a system.

Outcomes

- Reset the lost root user password.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command runs a start script that determines whether the `servera` machine is reachable on the network. It also resets the root password to a random string and sets a higher time-out for the GRUB2 menu.

```
[student@workstation ~]$ lab start boot-resetting
```

Instructions

- 1. Reboot `servera`, and interrupt the countdown in the boot-loader menu.
 - 1.1. Locate the icon for the `servera` console, as appropriate for your classroom environment, and then open the console.
Send `Ctrl+Alt+Del` to your system by using the relevant button or menu entry.
 - 1.2. When the boot-loader menu appears, press any key to interrupt the countdown, except `Enter`.
- 2. Edit the rescue kernel boot-loader entry, in memory, to abort the boot process just after the kernel mounts all the file systems, but before it hands over control to `systemd`.
 - 2.1. Use the cursor keys to highlight the rescue kernel entry (the one with the *rescue* word in its name).
 - 2.2. Press `e` to edit the current entry.
 - 2.3. Use the cursor keys to navigate to the line that starts with `linux`.
 - 2.4. Remove any `console=` options from the line.



Important

If you do not remove the `console=` options from the GRUB entry for the rescue kernel, the root prompt might be started on the wrong console and you might not be able to access it easily.

- 2.5. Press `End` to move the cursor to the end of the line.

2.6. Append a space followed by the `rd.break` option to the end of the line.



Note

If it is difficult for you to see the text in the console, then consider changing the resolution when editing the kernel line in the boot loader entry.

To change the console resolution, add either `video=640x480` or `vga=ask` on the line that starts with the `linux` word, after `rd.break`. For most consoles, a resolution of 640x480 is enough. By using `vga=ask`, you can choose a more suitable resolution for your environment.

2.7. Press `Ctrl+x` to boot by using the modified configuration.

- ▶ 3. Press `Enter` to perform maintenance. At the `sh-5.1#` prompt, remount the `/sysroot` file-system as read/write, and then use the `chroot` command to enter a `chroot` jail at `/sysroot`.

```
switch_root:/# mount -o remount,rw /sysroot
...output omitted...
switch_root:/# chroot /sysroot
sh-5.1#
```

- ▶ 4. Change the root password back to `redhat`.

```
sh-5.1# passwd root
Changing password for user root.
New password: redhat
BAD PASSWORD: The password is shorter than 8 characters
Retype new password: redhat
passwd: all authentication tokens updated successfully.
```

- ▶ 5. Configure the system to automatically perform a full SELinux relabeling after booting. This step is necessary because the `passwd` command re-creates the `/etc/shadow` file without an SELinux context.

```
sh-5.1# touch /.autorelabel
```

- ▶ 6. Type `exit` twice to continue booting your system as usual. The system runs an SELinux relabel operation, and then reboots automatically. When the system is up, verify your work by logging in as `root` at the console.

Finish

On the workstation machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish boot-resetting
```

Repair File-system Issues at Boot

Objectives

- Manually repair file-system configuration or corruption issues that stop the boot process.

File-system Issues

During the boot process, the `systemd` service mounts the persistent file systems that are defined in the `/etc/fstab` file.

Errors in the `/etc/fstab` file or corrupted file systems can block a system from completing the boot process. In some failure scenarios, the system breaks out of the boot process and opens an emergency shell that requires the `root` user password.

The following list describes some common file-system mounting issues when parsing the `/etc/fstab` file during the boot process:

Corrupted file system

The `systemd` service attempts to repair the file system. If the problem cannot be automatically repaired, then the system opens an emergency shell.

Nonexistent device or UUID

The `systemd` service times out waiting for the device to become available. If the device does not respond, then the system opens an emergency shell.



Note

If the mount point is not present, then Red Hat Enterprise Linux 9 automatically creates it during the boot process.

Repair File-system Issues at Boot

To access a system that cannot complete booting because of file-system issues, the `systemd` architecture provides an `emergency` boot target, which opens an emergency shell that requires the `root` password for access.

The next example demonstrates the boot process output when the system finds a file-system issue and switches to the `emergency` target:

```
...output omitted...
[*] A start job is running for /dev/vda2 (27s / 1min 30s)
[ TIME ] Timed out waiting for device /dev/vda2.
[DEPEND] Dependency failed for /mnt/mountfolder
[DEPEND] Dependency failed for Local File Systems.
[DEPEND] Dependency failed for Mark need to relabel after reboot.
...output omitted...
[ OK ] Started Emergency Shell.
[ OK ] Reached target Emergency Mode.
```

```
...output omitted...
Give root password for maintenance
(or press Control-D to continue):
```

The `systemd` daemon failed to mount the `/dev/vda2` device and timed out. Because the device is not available, the system opens an emergency shell for maintenance access.

To repair file-system issues when your system opens an emergency shell, first locate the errant file system, and then find and repair the fault. Now reload the `systemd` configuration to retry the automatic mounting.

Use the `mount` command to find which file systems are currently mounted by the `systemd` daemon.

```
[root@host ~]# mount
...output omitted...
/dev/vda1 on / type xfs (ro,relatime,seclabel,attr2,inode64,noquota)
...output omitted...
```

If the root file system is mounted with the `ro` (read-only) option, then you cannot edit the `/etc/fstab` file. Temporarily remount the root file system with the `rw` (read/write) option, if necessary, before opening the `/etc/fstab` file. With the `remount` option, an in-use file system can change its mount parameters without unmounting the file system.

```
[root@host ~]# mount -o remount,rw /
```

Try to mount all the file systems that are listed in the `/etc/fstab` file by using the `mount --all` option. This option mounts processes on every file-system entry, but skips those file systems that are already mounted. The command displays any errors that occur when mounting a file system.

```
[root@host ~]# mount --all
mount: /mnt/mountfolder: mount point does not exist.
```

In this scenario, where the `/mnt/mountfolder` mount directory does not exist, create the `/mnt/mountfolder` directory before reattempting the mount. Other error messages can occur, including typing errors in the entries, or wrong device names or UUIDs.

After you corrected all issues in the `/etc/fstab` file, inform the `systemd` daemon to register the new `/etc/fstab` file by using the `systemctl daemon-reload` command. Then, reattempt mounting all the entries.

```
[root@host ~]# systemctl daemon-reload
[root@host ~]# mount --all
```



Note

The `systemd` service processes the `/etc/fstab` file by transforming each entry into a `.mount` type `systemd` unit configuration and then starting the unit as a service. The `daemon-reload` option requests the `systemd` daemon to rebuild and reload all unit configurations.

If the `mount --all` command succeeds without further errors, then the final test is to verify that file-system mounting is successful during a system boot. Reboot the system and wait for the boot to complete normally.

```
[root@host ~]# systemctl reboot
```

For quick testing in the `/etc/fstab` file, use the `nofail` mount entry option. Using the `nofail` option in an `/etc/fstab` entry enables the system to boot even if that file-system mount is unsuccessful. This option must not be used with production file systems that must always mount. With the `nofail` option, an application could start when its file-system data is missing, with possibly severe consequences.



References

`systemd-fsck(8)`, `systemd-fstab-generator(8)`, and `systemd.mount(5)`
man pages

► Guided Exercise

Repair File-system Issues at Boot

Recover a system from a misconfiguration in the `/etc/fstab` file where the boot process fails.

Outcomes

- Diagnose `/etc/fstab` file issues and use emergency mode to recover the system.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start boot-repairing
```

Instructions

- 1. Access the `servera` machine console and notice that the boot process is stuck early on.
 - 1.1. Locate the icon for the `servera` console, as appropriate for your classroom environment. Open the console.
Notice that a start job does not seem to complete. Consider a possible cause for this behavior.
 - 1.2. Reboot the `servera` machine, by sending `Ctrl+Alt+Del` to your system by using the relevant button or menu entry. With this boot problem, this key sequence might not immediately abort the running job, and you might have to wait for it to time out before the system reboots.
If you wait for the task to time out without sending `Ctrl+Alt+Del`, then the system eventually spawns an emergency shell by itself.
 - 1.3. When the boot-loader menu appears, press any key to interrupt the countdown, except the `Enter` key.
- 2. Looking at the error from the previous boot, parts of the system still seem to be functioning. Use `redhat` as the `root` user password to try an emergency boot.
 - 2.1. Use the cursor keys to highlight the default boot loader entry.
 - 2.2. Press the `e` key to edit the current entry.
 - 2.3. Use the cursor keys to navigate to the line that starts with the `linux` word.
 - 2.4. Press `End` to move the cursor to the end of the line.
 - 2.5. Append the `systemd.unit=emergency.target` string to the end of the line.

**Note**

If it is difficult for you to see the text in the console, consider changing the resolution when editing the kernel line in the boot loader entry.

To change the console resolution, add either `video=640x480` or `vga=ask` on the line that starts with the `linux` word, after `systemd.unit=emergency.target`. For most consoles, a resolution of 640x480 is enough. By using `vga=ask`, you can choose a more suitable resolution for your environment.

2.6. Press `Ctrl+x` to boot by using the modified configuration.

- ▶ 3. Log in to emergency mode.

```
Give root password for maintenance
(or press Control-D to continue): redhat
[root@servera ~]#
```

- ▶ 4. Determine which file systems the `systemd` daemon currently mounts. The `systemd` daemon mounts the root file system in read-only mode.

```
[root@servera ~]# mount
...output omitted...
/dev/vda4 on / type xfs
(ro,relatime,seclabel,attr2,inode64,logbufs=8,logbsize=32k,noquota)
...output omitted...
```

- ▶ 5. Remount the root file system in read/write mode.

```
[root@servera ~]# mount -o remount,rw /
```

- ▶ 6. Try to mount all the other file systems. The `--all` (`-a`) option mounts all the listed file systems in the `/etc/fstab` file that are not yet mounted.

```
[root@servera ~]# mount -a
mount: /RemoveMe: mount point does not exist.
```

- ▶ 7. Edit the `/etc/fstab` file to fix the issue.

7.1. Remove or comment out the incorrect line from the `/etc/fstab` file.

```
[root@servera ~]# cat /etc/fstab
...output omitted...
# /dev/sdz1 /RemoveMe xfs defaults 0 0
```

7.2. Reload the `systemd` daemon for the system to register the new `/etc/fstab` file configuration.

```
[root@servera ~]# systemctl daemon-reload
```

- ▶ 8. Verify that the `/etc/fstab` file is now correct by attempting to mount all entries.

```
[root@servera ~]# mount -a
```

- ▶ 9. Reboot the system and wait for the boot to complete. The system should now boot normally.

```
[root@servera ~]# systemctl reboot
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish boot-repairing
```

► Lab

Control the Boot Process

Reset the `root` password on a system, recover from a misconfiguration, and set the default boot target.

Outcomes

- Reset a lost password for the `root` user.
- Diagnose and fix boot issues.
- Set the default `systemd` target.

Before You Begin

As the student user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start boot-review
```

Instructions

1. On the `serverb` machine, reset the password to `redhat` for the `root` user.
Locate the icon for the `serverb` machine console as appropriate for your classroom environment, and then open the console.
2. In the boot-loader menu, select the default kernel boot-loader entry. The system fails to boot, because a start job does not complete successfully. Fix the issue from the console of the `serverb` machine.
3. Change the default `systemd` target on the `serverb` machine for the system to automatically start a graphical interface when it boots.
No graphical interface is installed on the `serverb` machine. Set only the default target, and do not install the packages.

Evaluation

As the student user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade boot-review
```


Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish boot-review
```

► Solution

Control the Boot Process

Reset the `root` password on a system, recover from a misconfiguration, and set the default boot target.

Outcomes

- Reset a lost password for the `root` user.
- Diagnose and fix boot issues.
- Set the default `systemd` target.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start boot-review
```

Instructions

1. On the `serverb` machine, reset the password to `redhat` for the `root` user.
Locate the icon for the `serverb` machine console as appropriate for your classroom environment, and then open the console.
 - 1.1. Send `Ctrl+Alt+Del` to your system by using the relevant button or menu entry.
 - 1.2. When the boot-loader menu appears, press any key to interrupt the countdown, except the `Enter` key.
 - 1.3. Use the cursor keys to highlight the rescue kernel boot-loader entry (the one with the *rescue* word in its name).
 - 1.4. Press `e` to edit the current entry.
 - 1.5. Use the cursor keys to navigate the line that starts with the `linux` text.
 - 1.6. Remove all `console=` options from the rescue kernel entry.



Important

If you do not remove the `console=` options from the GRUB entry for the rescue kernel, the root prompt might be started on the wrong console and you might not be able to access it easily.

- 1.7. Press `Ctrl+e` to move the cursor to the end of the line.

- 1.8. Append a space followed by the `rd.break` option to the end of the line.

**Note**

If it is difficult for you to see the text in the console, consider changing the resolution when editing the kernel line in the boot loader entry.

To change the console resolution, add either `video=640x480` or `vga=ask` on the line that starts with the `linux` word, after `rd.break`. For most consoles, a resolution of 640x480 is enough. By using `vga=ask`, you can choose a more suitable resolution for your environment.

- 1.9. Press `Ctrl+x` to boot by using the modified configuration.
- 1.10. Press `Enter` to enter the maintenance mode.
- 1.11. At the `switch_root:/#` prompt, remount the `/sysroot` file system as writable, and then use the `chroot` command for the `/sysroot` directory.

```
switch_root:/# mount -o remount,rw /sysroot
...output omitted...
switch_root:/# chroot /sysroot
sh-5.1#
```

- 1.12. Set `redhat` as the password for the `root` user.

```
sh-5.1# passwd root
Changing password for user root.
New password: redhat
BAD PASSWORD: The password is shorter than 8 characters
Retype new password: redhat
passwd: all authentication tokens updated successfully.
```

- 1.13. Configure the system to perform a full SELinux relabeling after booting.

```
sh-5.1# touch /.autorelabel
```

- 1.14. Exit the `chroot` environment and the `switch_root:/#` prompt. After the file system is relabeled, the system prompts to enter maintenance mode. However, if you wait, then it completes the reboot and shows the boot-loader menu.
2. In the boot-loader menu, select the default kernel boot-loader entry. The system fails to boot, because a start job does not complete successfully. Fix the issue from the console of the `serverb` machine.
 - 2.1. Boot the system into emergency mode. Reboot the `serverb` machine by sending `Ctrl+Alt+Del` to your system by using the relevant button or menu entry.
 - 2.2. When the boot-loader menu appears, press any key to interrupt the countdown, except `Enter`.
 - 2.3. Use the cursor keys to highlight the default boot-loader entry.
 - 2.4. Press `e` to edit the current entry.

- 2.5. Use the cursor keys to navigate the line that starts with the `linux` text.
- 2.6. Press `Ctrl+E` to move the cursor to the end of the line.
- 2.7. Append the `systemd.unit=emergency.target` text to the end of the line.
- 2.8. Press `Ctrl+X` to boot by using the modified configuration.
- 2.9. Log in to emergency mode.

```
Give root password for maintenance
(or press Control-D to continue): redhat
[root@serverb ~]#
```

- 2.10. Remount the `/` file system as writable.

```
[root@serverb ~]# mount -o remount,rw /
...output omitted...
```

- 2.11. Mount all file systems.

```
[root@serverb ~]# mount -a
mount: /olddata: can't find UUID=4d5c85a5-8921-4a06-8aff-80567e9689bc.
```

- 2.12. Edit the `/etc/fstab` file to remove or comment out the incorrect line that mounts the `/olddata` mount point.

```
[root@serverb ~]# vim /etc/fstab
...output omitted...
#UUID=4d5c85a5-8921-4a06-8aff-80567e9689bc /olddata xfs defaults 0 0
```

- 2.13. Update the `systemd` daemon for the system to register the changes in the `/etc/fstab` file configuration.

```
[root@serverb ~]# systemctl daemon-reload
```

- 2.14. Verify that the `/etc/fstab` file configuration is correct by attempting to mount all entries.

```
[root@serverb ~]# mount -a
```

- 2.15. Reboot the system and wait for the boot to complete. The system should now boot normally.

```
[root@serverb ~]# systemctl reboot
```

3. Change the default `systemd` target on the `serverb` machine for the system to automatically start a graphical interface when it boots.
No graphical interface is installed on the `serverb` machine. Set only the default target, and do not install the packages.

3.1. Log in to the `serverb` machine as the `student` user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
[student@serverb ~]$ sudo -i
[sudo] password for student: student
[root@serverb ~]#
```

3.2. Set the `graphical.target` as the default target.

```
[root@serverb ~]# systemctl set-default graphical.target
Removed /etc/systemd/system/default.target.
Created symlink /etc/systemd/system/default.target → /usr/lib/systemd/system/
graphical.target.
```

3.3. Verify that the correct default is set.

```
[root@serverb ~]# systemctl get-default
graphical.target
```

3.4. Return to the `workstation` machine as the `student` user.

```
[root@serverb ~]# exit
logout
[student@serverb ~]$ exit
logout
Connection to serverb closed.
```

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade boot-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish boot-review
```

Summary

- The `systemctl reboot` and `systemctl poweroff` commands reboot and power down a system, respectively.
- The `systemctl isolate target-name.target` command switches to a new target at runtime.
- The `systemctl get-default` and `systemctl set-default` commands can query and set the default target.
- You can use the `rd.break` option on the kernel command line to interrupt the boot process before control is handed over from the `initramfs` image. The root file system is mounted read-only under `/sysroot`.
- The emergency target can diagnose and fix file-system issues.

Chapter 11

Manage Network Security

Goal

Control network connections to services with the system firewall and SELinux rules.

Objectives

- Accept or reject network connections to system services with `firewalld` rules.
- Verify that network ports have the correct SELinux type for services to bind to them.

Sections

- Manage Server Firewalls (and Guided Exercise)
- Control SELinux Port Labeling (and Guided Exercise)

Lab

- Manage Network Security

Manage Server Firewalls

Objectives

- Accept or reject network connections to system services with `firewalld` rules.

Firewall Architecture Concepts

The Linux kernel provides the `netfilter` framework for network traffic operations such as packet filtering, network address translation, and port translation. The `netfilter` framework includes *hooks* for kernel modules to interact with network packets as they traverse a system's network stack. Fundamentally, `netfilter` hooks are kernel routines that intercept events (for example, a packet that enters an interface) and run other related routines (for example, firewall rules).

The nftables Framework

The `nftables` packet classification framework builds on the `netfilter` framework to apply firewall rules to network traffic. In Red Hat Enterprise Linux 9, the `nftables` framework is the system firewall core, and it replaces the deprecated `iptables` framework.

The `nftables` framework provides many advantages over `iptables`, including improved usability and more efficient rule sets. For example, the `iptables` framework required a rule for each protocol, but `nftables` rules can apply to both IPv4 and IPv6 traffic simultaneously. The `iptables` framework required using different tools, such as `iptables`, `ip6tables`, `arptables`, and `ebtables`, for each protocol. By contrast, the `nftables` framework uses the single `nft` user-space utility to manage all protocols through a single interface.



Note

Convert earlier `iptables` configuration files into their `nftables` equivalents by using the `iptables-translate` and `ip6tables-translate` utilities.

The firewalld Service

The `firewalld` service is a dynamic firewall manager, and is the recommended front end to the `nftables` framework. The Red Hat Enterprise Linux 9 distribution includes the `firewalld` package.

The `firewalld` service simplifies firewall management by classifying network traffic into *zones*. A network packet's assigned zone depends on criteria such as the source IP address of the packet or the incoming network interface. Each zone has its own list of ports and services that are either open or closed.

**Note**

For laptops or other machines that often change networks, the `NetworkManager` service can automatically set the firewall zone for a connection. This service is useful when switching between home, work, and public wireless networks. A user might want their system's `sshd` service to be reachable when connected to their home or corporate networks, but not when connected to a public wireless network in the local coffee shop.

The `firewalld` service ensures the source address for every incoming packet into the system. If that source address is assigned to a specific zone, then the rules for that zone apply. If the source address is not assigned to a zone, then the `firewalld` service associates the packet with the zone for the incoming network interface, and the rules for that zone apply. If the network interface is not associated with a zone, then the `firewalld` service sends the packet to the default zone.

The default zone is not a separate zone but rather an assigned designation to an existing zone. Initially, the `firewalld` service designates the `public` zone as default, and maps the `lo` loopback interface to the `trusted` zone.

Most zones allow traffic through the firewall if it matches a list of particular ports and protocols, such as `631/udp`, or a predefined service configuration, such as `ssh`. Normally, if the traffic does not match a permitted port and protocol or service, then it is rejected. The `trusted` zone, which permits all traffic by default, is an exception.

Predefined Zones

The `firewalld` service uses predefined zones, which you can customize. By default, all zones allow any incoming traffic that is part of an existing session that system initiated, and also allow all outgoing traffic. The following table details the initial zone configuration.

Default Configuration of Firewalld Zones

Zone name	Default configuration
<code>trusted</code>	Allow all incoming traffic.
<code>home</code>	Reject incoming traffic unless related to outgoing traffic or matching the <code>ssh</code> , <code>mdns</code> , <code>ipp-client</code> , <code>samba-client</code> , or <code>dhcpv6-client</code> predefined services.
<code>internal</code>	Reject incoming traffic unless related to outgoing traffic or matching the <code>ssh</code> , <code>mdns</code> , <code>ipp-client</code> , <code>samba-client</code> , or <code>dhcpv6-client</code> predefined services (same as the <code>home</code> zone to start with).
<code>work</code>	Reject incoming traffic unless related to outgoing traffic or matching the <code>ssh</code> , <code>ipp-client</code> , or <code>dhcpv6-client</code> predefined services.
<code>public</code>	Reject incoming traffic unless related to outgoing traffic or matching the <code>ssh</code> or <code>dhcpv6-client</code> predefined services. <i>The default zone for newly added network interfaces.</i>

Zone name	Default configuration
external	Reject incoming traffic unless related to outgoing traffic or matching the ssh predefined service. Outgoing IPv4 traffic that is forwarded through this zone is <i>masqueraded</i> to appear that it originated from the IPv4 address of the outgoing network interface.
dmz	Reject incoming traffic unless related to outgoing traffic or matching the ssh predefined service.
block	Reject all incoming traffic unless related to outgoing traffic.
drop	Drop all incoming traffic unless related to outgoing traffic (do not even respond with ICMP errors).

For a list of available predefined zones and their intended use, see the `firewalld.zones(5)` man page.

Predefined Services

The `firewalld` service includes predefined configurations for common services, to simplify setting firewall rules. For example, instead of researching the relevant ports for an NFS server, use the predefined `nfs` configuration create rules for the correct ports and protocols. The following table lists some predefined service configurations that might be active in your default `firewalld` zone.

Selected Predefined Firewalld Services

Service name	Configuration
ssh	Local SSH server. Traffic to 22/tcp.
dhcpv6-client	Local DHCPv6 client. Traffic to 546/udp on the fe80::/64 IPv6 network.
ipp-client	Local IPP printing. Traffic to 631/udp.
samba-client	Local Windows file and print sharing client. Traffic to 137/udp and 138/udp.
mdns	Multicast DNS (mDNS) local-link name resolution. Traffic to 5353/udp to the 224.0.0.251 (IPv4) or ff02::fb (IPv6) multicast addresses.
cockpit	Red Hat Enterprise Linux web-based interface for managing and monitoring your local and remote system. Traffic to 9090 port.

The `firewalld` package includes many predefined service configurations. You can list the services with the `firewall-cmd --get-services` command.

```
[root@host ~]# firewall-cmd --get-services
RH-Satellite-6 RH-Satellite-6-capsule amanda-client amanda-k5-client amqp amqps
apcupsd audit bacula bacula-client bb bgp bitcoin bitcoin-rpc bitcoin-testnet
bitcoin-testnet-rpc bittorrent-lsd ceph ceph-mon cfengine cockpit collectd
...output omitted...
```

If the predefined service configurations are not appropriate for your scenario, then you can manually specify the required ports and protocols. You can use the web console graphical interface to review predefined services and manually define more ports and protocols.

Configure the firewalld Daemon

The following list shows two common ways that system administrators use to interact with the `firewalld` service:

- The web console graphical interface
- The `firewall-cmd` command-line tool

Configure Firewall Services with the Web Console

To manage firewall services with the web console, you must log in and escalate privileges. You can escalate privileges by clicking the **Limited access** or **Turn on administrative access** buttons. Then, enter your password when prompted. The administrative mode elevates privileges based on your user's `sudo` configuration. As a security reminder, remember to toggle back to limited access mode after you perform the system task that requires administrative privileges.

Click the **Networking** option in the left navigation menu to display the **Firewall** section in the main networking page. Click the **Edit rules and zones** button zones to navigate to the **Firewall** page.




Figure 11.1: The web console networking page

The **Firewall** page displays active zones and their allowed services. Click the arrow (>) button to the left of a service name to view its details. To add a service to a zone, click the **Add services** button in the upper right corner of the applicable zone.




Figure 11.2: The web console firewall page

The **Add Services** page displays the available predefined services.

Add services to public zone [X]

☒ Services ☐ Custom ports

Filter services

- ☐ RH-Satellite-6
TCP: 5000, 5646-5647, 5671, 8000, 8080, 9090, 8140 UDP: 68
- ☐ RH-Satellite-6-capsule
TCP: 8443, 5000, 5646-5647, 5671, 8000, 8080, 9090
- ☐ amanda-client
TCP: 10080 UDP: 10080
- ☐ amanda-k5-client

Add services Cancel

Figure 11.3: The web console add services menu

To select a service, scroll through the list or enter a selection in the **Filter services** text box. In the following example, the `http` string filters the options to web-related services. Select the checkbox to the left of the service to allow it through the firewall. Click the **Add services** button to complete the process.

Add services to public zone [X]

☒ Services ☐ Custom ports

Filter services

- ☒ http
TCP: 80
- ☐ https
TCP: 443
- ☐ wbem-http
TCP: 5988
- ☐ wbem-https

Add services Cancel

Figure 11.4: The web console add services menu options

The interface returns to the **Firewall** page, where you can review the updated allowed services list.




Figure 11.5: The web console firewall overview

Configure the Firewall from the Command Line

The `firewall-cmd` command interfaces with the `firewalld` daemon. It is installed as part of the `firewalld` package, and is available for administrators who prefer to work on the command line, for working on systems without a graphical environment, or for scripting a firewall setup.

The following table lists often use `firewall-cmd` commands, along with an explanation. Most commands work on the *runtime* configuration, unless the `--permanent` option is specified. If the `--permanent` option is specified, then you must activate the setting by also running the `firewall-cmd --reload` command, which reads the current permanent configuration and applies it as the new runtime configuration. Many of the listed commands take the `--zone=ZONE` option to find which zone they affect. Where a netmask is required, use CIDR notation, such as `192.168.1/24`.

firewall-cmd commands	Explanation
<code>--get-default-zone</code>	Query the current default zone.
<code>--set-default-zone=ZONE</code>	Set the default zone. This default zone changes both the runtime and the permanent configuration.
<code>--get-zones</code>	List all available zones.
<code>--get-active-zones</code>	List all zones that are currently in use (with an interface or source that is tied to them), along with their interface and source information.
<code>--add-source=CIDR [--zone=ZONE]</code>	Route all traffic from the IP address or network/netmask to the specified zone. If no <code>--zone=</code> option is provided, then the default zone is used.
<code>--remove-source=CIDR [--zone=ZONE]</code>	Remove the rule that routes all traffic from the zone that comes from the IP address or network. If no <code>--zone=</code> option is provided, then the default zone is used.

firewall-cmd commands	Explanation
<code>--add-interface=INTERFACE</code> <code>[--zone=ZONE]</code>	Route all traffic from <i>INTERFACE</i> to the specified zone. If no <code>--zone=</code> option is provided, then the default zone is used.
<code>--change-interface=INTERFACE</code> <code>[--zone=ZONE]</code>	Associate the interface with <i>ZONE</i> instead of its current zone. If no <code>--zone=</code> option is provided, then the default zone is used.
<code>--list-all [--zone=ZONE]</code>	List all configured interfaces, sources, services, and ports for <i>ZONE</i> . If no <code>--zone=</code> option is provided, then the default zone is used.
<code>--list-all-zones</code>	Retrieve all information for all zones (interfaces, sources, ports, and services).
<code>--add-service=SERVICE</code> <code>[--zone=ZONE]</code>	Allow traffic to <i>SERVICE</i> . If no <code>--zone=</code> option is provided, then the default zone is used.
<code>--add-port=PORT/PROTOCOL</code> <code>[--zone=ZONE]</code>	Allow traffic to the <i>PORT/PROTOCOL</i> ports. If no <code>--zone=</code> option is provided, then the default zone is used.
<code>--remove-service=SERVICE</code> <code>[--zone=ZONE]</code>	Remove <i>SERVICE</i> from the allowed list for the zone. If no <code>--zone=</code> option is provided, then the default zone is used.
<code>--remove-port=PORT/PROTOCOL</code> <code>[--zone=ZONE]</code>	Remove the <i>PORT/PROTOCOL</i> ports from the allowed list for the zone. If no <code>--zone=</code> option is provided, then the default zone is used.
<code>--reload</code>	Drop the runtime configuration and apply the persistent configuration.

The following example sets the default zone to `dmz`, assigns all traffic coming from the `192.168.0.0/24` network to the `internal` zone, and opens the network ports for the `mysql` service on the `internal` zone.

```
[root@host ~]# firewall-cmd --set-default-zone=dmz
[root@host ~]# firewall-cmd --permanent --zone=internal \
--add-source=192.168.0.0/24
[root@host ~]# firewall-cmd --permanent --zone=internal --add-service=mysql
[root@host ~]# firewall-cmd --reload
```

As another example, to add all the incoming traffic from the `172.25.25.11` single IPv4 address to the `public` zone, use the following commands:

```
[root@host ~]# firewall-cmd --permanent --zone=public \
--add-source=172.25.25.11/32
[root@host ~]# firewall-cmd --reload
```

**Note**

For situations where the basic syntax is not enough, you can add *rich-rules* to write complex rules. If even the rich-rules syntax is not enough, then you can also use Direct Configuration rules (which use raw `nft` syntax mixed in with `firewalld` rules). These advanced configurations are beyond the scope of this chapter.

**References**

`firewall-cmd(1)`, `firewalld(1)`, `firewalld.zone(5)`, `firewalld.zones(5)`, and `nft(8)` man pages

► Guided Exercise

Manage Server Firewalls

Control access to system services by adjusting system firewall rules with the `firewalld` service.

Outcomes

- Configure firewall rules to control access to services.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start netsecurity-firewalls
```

Instructions

- 1. Log in to the `servera` machine as the `student` user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

- 2. Install the `httpd` and `mod_ssl` packages. These packages provide the Apache web server and the necessary extensions for the web server to serve content over SSL.

```
[root@servera ~]# dnf install httpd mod_ssl
...output omitted...
Is this ok [y/N]: y
...output omitted...
Complete!
```

- 3. Create the `/var/www/html/index.html` file. Add one line of text that reads: `I am servera.`

```
[root@servera ~]# echo 'I am servera.' > /var/www/html/index.html
```

- 4. Start and enable the `httpd` service.

```
[root@servera ~]# systemctl enable --now httpd
Created symlink /etc/systemd/system/multi-user.target.wants/httpd.service → /usr/lib/systemd/system/httpd.service.
```

- 5. Return to the workstation machine as the student user.

```
[root@servera ~]# exit
logout
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

- 6. From workstation, try to access the web server on servera by using both the 80/TCP clear-text port and the 443/TCP SSL encapsulated port. Both attempts should fail.

6.1. The `curl` command should fail.

```
[student@workstation ~]$ curl http://servera.lab.example.com
curl: (7) Failed to connect to servera.lab.example.com port 80: No route to host
```

6.2. The `curl` command with the `-k` option for insecure connections should also fail.

```
[student@workstation ~]$ curl -k https://servera.lab.example.com
curl: (7) Failed to connect to servera.lab.example.com port 443: No route to host
```

- 7. Verify that the `firewalld` service on servera is enabled and running.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]# systemctl status firewalld
● firewalld.service - firewalld - dynamic firewall daemon
   Loaded: loaded (/usr/lib/systemd/system/firewalld.service; enabled; vendor preset: enabled)
   Active: active (running) since Wed 2022-04-13 11:22:50 EDT; 7min ago
     Docs: man:firewalld(1)
  Main PID: 768 (firewalld)
    Tasks: 2 (limit: 10798)
   Memory: 39.9M
      CPU: 584ms
   CGroup: /system.slice/firewalld.service
           └─768 /usr/bin/python3 -s /usr/sbin/firewalld --nofork --nopid

Apr 13 11:22:49 servera.lab.example.com systemd[1]: Starting firewalld - dynamic firewall daemon...
Apr 13 11:22:50 servera.lab.example.com systemd[1]: Started firewalld - dynamic firewall daemon.
```

- ▶ 8. Add the `https` service to the `public` firewall zone.

8.1. Verify that the default firewall zone is set to the `public` zone.

```
[root@servera ~]# firewall-cmd --get-default-zone
public
```

8.2. If the earlier step does not return `public` as the default zone, then correct it with the following command:

```
[root@servera ~]# firewall-cmd --set-default-zone public
```

8.3. Add the `https` service to the permanent configuration for the `public` network zone. Confirm your configuration.

```
[root@servera ~]# firewall-cmd --permanent --add-service=https
success
[root@servera ~]# firewall-cmd --reload
success
[root@servera ~]# firewall-cmd --permanent --zone=public --list-all
public
  target: default
  icmp-block-inversion: no
  interfaces:
  sources:
  services: cockpit dhcpv6-client https ssh
  ports:
  protocols:
  forward: yes
  masquerade: no
  forward-ports:
  source-ports:
  icmp-blocks:
  rich rules:
```

- ▶ 9. From `workstation`, open Firefox and log in to the web console that is running on `servera` to verify the `https` service to the `public` firewall zone.
 - 9.1. Open Firefox and navigate to `https://servera.lab.example.com:9090` to access the web console. Click **Advanced** and **Accept the Risk and Continue** to accept the self-signed certificate.
 - 9.2. Log in as the `student` user with `student` as the password.
 - 9.3. Click **Turn on administrative access** and enter the `student` password again.
 - 9.4. Click **Networking** in the left navigation bar.
 - 9.5. Click **Edit rules and zones** in the **Firewall** section of the **Networking** page.
 - 9.6. Verify that the `https` service is listed in the **Service** column.
- ▶ 10. Return to a terminal on `workstation`, and verify your work by attempting to access the `servera` web server.

10.1. Return to the `workstation` machine as the `student` user.

```
[root@servera ~]# exit
logout
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

10.2. Verify the access to the `http://servera.lab.example.com` web server.

```
[student@workstation ~]$ curl http://servera.lab.example.com
curl: (7) Failed to connect to servera.lab.example.com port 80: No route to host
```

10.3. Verify the access to the `http://servera.lab.example.com` web server through the port 443 for insecure connection.

```
[student@workstation ~]$ curl -k https://servera.lab.example.com
I am servera.
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from earlier exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish netsecurity-firewalls
```

Control SELinux Port Labeling

Objectives

- Verify that network ports have the correct SELinux type for services to bind to them.

SELinux Port Labeling

In addition to file context and process type labeling, SELinux labels network ports with an SELinux context. SELinux controls network access by labeling the network ports and including rules in a service's targeted policy. For example, the SSH targeted policy includes the 22/TCP port with an `ssh_port_t` port context label. In the HTTP policy, the default 80/TCP and 443/TCP ports use an `http_port_t` port context label.

When a targeted process attempts to open a port for listening, SELinux verifies that the policy includes entries that enable the binding of the process and the context. SELinux can then block a rogue service from taking over ports that other legitimate network services use.

Manage SELinux Port Labeling

If a service attempts to listen on a nonstandard port, and the port is not labeled with the correct SELinux type, then SELinux might block the attempt. You can correct this problem by changing the SELinux context on the port.

Typically, the targeted policy already labeled all expected ports with the correct type. For example, because port 8008/TCP is often used for web applications, that port is already labeled with `http_port_t`, which is the default port type that a web server uses. Individual ports can be labeled with only one port context.

List Port Labels

Use the `grep` command to filter the port number.

```
[root@host ~]# grep gopher /etc/services
gopher          70/tcp          # Internet Gopher
gopher          70/udp
```

Use the `semanage` command to list the current port label assignments.

```
[root@host ~]# semanage port -l
...output omitted...
http_cache_port_t      tcp    8080, 8118, 8123, 10001-10010
http_cache_port_t      udp    3130
http_port_t            tcp    80, 81, 443, 488, 8008, 8009, 8443, 9000
...output omitted...
```

Use the `grep` command to filter the SELinux port label by using the service name.

```
[root@host ~]# semanage port -l | grep ftp
ftp_data_port_t      tcp      20
ftp_port_t           tcp      21, 989, 990
ftp_port_t           udp      989, 990
tftp_port_t          udp      69
```

A port label can appear in the list many times for each supported networking protocol.

Use the `grep` command to filter the SELinux port label by using the port number.

```
[root@host ~]# semanage port -l | grep -w 70
gopher_port_t        tcp      70
gopher_port_t        udp      70
```

Manage Port Bindings

Use the `semanage` command to assign new port labels, remove port labels, and modify existing ones.



Important

Almost all of the services that are included in the RHEL distribution provide an SELinux policy module, which includes that service's default port contexts. You cannot change default port labels by using the `semanage` command. Instead, you must modify and reload the targeted service's policy module. Writing and generating policy modules is not discussed in this course.

You can label a new port with an existing port context label (type). The `semanage port` command's `-a` option adds a new port label; the `-t` option denotes the type; and the `-p` option denotes the protocol.

```
[root@host ~]# semanage port -a -t port_label -p tcp|udp PORTNUMBER
```

In the following example, enable the `gopher` service to listen on the 71/TCP port:

```
[root@host~]# semanage port -a -t gopher_port_t -p tcp 71
```

To view local changes to the default policy, use the `semanage port` command's `-C` option.

```
[root@host~]# semanage port -l -C
SELinux Port Type      Proto  Port Number
gopher_port_t          tcp    71
```

The targeted policies include many port types.

Service-specific SELinux man pages are named by using the service name plus `_selinux`. These man pages include service-specific information on SELinux types, Booleans, and port types, and are not installed by default. To view a list of all of the available SELinux man pages, install the package and then run a `man -k` keyword search for the `_selinux` string.

```
[root@host ~]# dnf -y install selinux-policy-doc
[root@host ~]# man -k _selinux
```

Use the `semanage` command for deleting a port label, with the `-d` option. In the following example, remove the binding of port 71/TCP to the `gopher_port_t` type:

```
[root@host ~]# semanage port -d -t gopher_port_t -p tcp 71
```

To change a port binding, when requirements change, use the `-m` option. This option is more efficient than deleting the earlier binding and adding the latest one.

For example, to modify port 71/TCP from `gopher_port_t` to `http_port_t`, use the following command:

```
[root@server ~]# semanage port -m -t http_port_t -p tcp 71
```

View the modification by using the `semanage` command.

```
[root@server ~]# semanage port -l -C
SELinux Port Type          Proto    Port Number

http_port_t                tcp      71
[root@server ~]# semanage port -l | grep http
http_cache_port_t          tcp      8080, 8118, 8123, 10001-10010
http_cache_port_t          udp      3130
http_port_t                tcp      71, 80, 81, 443, 488, 8008, 8009, 8443,
9000
pegasus_http_port_t         tcp      5988
pegasus_https_port_t        tcp      5989
```



References

`semanage(8)`, `semanage-port(8)`, and `*_selinux(8)` man pages

► Guided Exercise

Control SELinux Port Labeling

Configure your system to allow HTTP access on a nonstandard port.

Outcomes

- Configure a web server that is running on `servera` to successfully serve content that uses a nonstandard port.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command determines whether the `servera` machine is reachable on the network, installs the `httpd` service, and configures the firewall on `servera` to allow HTTP connections.

```
[student@workstation ~]$ lab start netsecurity-ports
```

Instructions

Your organization is deploying a new custom web application. The web application is running on a nonstandard port, in this case, 82/TCP.

A junior administrator already configured the application on your `servera` host. However, the web server content is not accessible.

- 1. Log in to `servera` as the `student` user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
[student@servera ~]$ sudo -i
[sudo] password for student: student
[root@servera ~]#
```

- 2. Try to fix the web content problem by restarting the `httpd` service.

2.1. Restart the `httpd.service`. This command is expected to fail.

```
[root@servera ~]# systemctl restart httpd.service
Job for httpd.service failed because the control process exited with error code.
See "systemctl status httpd.service" and "journalctl -xe" for details.
```

2.2. View the status of the `httpd` service. Note the `permission denied` error.


```
[root@servera ~]# systemctl status -l httpd.service
• httpd.service - The Apache HTTP Server
  Loaded: loaded (/usr/lib/systemd/system/httpd.service; enabled; vendor preset:
disabled)
  Active: failed (Result: exit-code) since Mon 2019-04-08 14:23:29 CEST; 3min 33s
ago
    Docs: man:httpd.service(8)
  Process: 28078 ExecStart=/usr/sbin/httpd $OPTIONS -DFOREGROUND (code=exited,
status=1/FAILURE)
 Main PID: 28078 (code=exited, status=1/FAILURE)
   Status: "Reading configuration..."

Apr 08 14:23:29 servera.lab.example.com systemd[1]: Starting The Apache HTTP
Server...
Apr 08 14:23:29 servera.lab.example.com httpd[28078]: (13)Permission denied:
AH00072: make_sock: could not bind to address [::]:82
Apr 08 14:23:29 servera.lab.example.com httpd[28078]: (13)Permission denied:
AH00072: make_sock: could not bind to address 0.0.0.0:82
Apr 08 14:23:29 servera.lab.example.com httpd[28078]: no listening sockets
available, shutting down
Apr 08 14:23:29 servera.lab.example.com httpd[28078]: AH00015: Unable to open logs
Apr 08 14:23:29 servera.lab.example.com systemd[1]: httpd.service: Main process
exited, code=exited, status=1/FAILURE
Apr 08 14:23:29 servera.lab.example.com systemd[1]: httpd.service: Failed with
result 'exit-code'.
Apr 08 14:23:29 servera.lab.example.com systemd[1]: Failed to start The Apache
HTTP Server.
```

2.3. Verify whether SELinux is blocking httpd from binding to the 82/TCP port.

```
[root@servera ~]# sealert -a /var/log/audit/audit.log
100% done
found 1 alerts in /var/log/audit/audit.log
-----

SELinux is preventing /usr/sbin/httpd from name_bind access on the tcp_socket port
82.

***** Plugin bind_ports (99.5 confidence) suggests *****

If you want to allow /usr/sbin/httpd to bind to network port 82
Then you need to modify the port type.
Do
# semanage port -a -t PORT_TYPE -p tcp 82 where PORT_TYPE is one of the following:
http_cache_port_t, http_port_t, jboss_management_port_t, jboss_messaging_port_t,
ntop_port_t, puppet_port_t.
...output omitted...
Raw Audit Messages
type=AVC msg=audit(1554726569.188:852): avc: denied { name_bind } for
pid=28393 comm="httpd" src=82 scontext=system_u:system_r:httpd_t:s0
tcontext=system_u:object_r:reserved_port_t:s0 tclass=tcp_socket permissive=0
...output omitted...
```

- ▶ **3.** Configure SELinux to allow the `httpd.service` to bind to the 82/TCP port, and then restart the `httpd.service` service.

3.1. Find an appropriate port type for the 82/TCP port.

The `http_port_t` type includes the default HTTP ports, 80/TCP and 443/TCP. This type is the correct port type for the web server.

```
[root@servera ~]# semanage port -l | grep http
http_cache_port_t      tcp      8080, 8118, 8123, 10001-10010
http_cache_port_t      udp      3130
http_port_t            tcp      80, 81, 443, 488, 8008, 8009, 8443, 9000
pegasus_http_port_t    tcp      5988
pegasus_https_port_t   tcp      5989
```

3.2. Assign the 82/TCP port the `http_port_t` type.

```
[root@servera ~]# semanage port -a -t http_port_t -p tcp 82
```

3.3. Restart the `httpd.service` service. This command should succeed.

```
[root@servera ~]# systemctl restart httpd.service
```

- ▶ **4.** Verify that you can now access the web server that runs on the 82/TCP port.

```
[root@servera ~]# curl http://servera.lab.example.com:82
Hello
```

- ▶ **5.** In a different terminal window, verify whether you can access the new web service from workstation.

```
[student@workstation ~]$ curl http://servera.lab.example.com:82
curl: (7) Failed to connect to servera.example.com:82; No route to host
```

That error means that you still cannot connect to the web service from `workstation`.

- ▶ **6.** On `servera`, open the 82/TCP port on the firewall.

6.1. Open the 82/TCP port in the permanent configuration, for the default zone on the firewall, on `servera`.

```
[root@servera ~]# firewall-cmd --permanent --add-port=82/tcp
success
```

6.2. Activate your firewall changes on `servera`.

```
[root@servera ~]# firewall-cmd --reload
success
```

- ▶ **7.** Access the web service from `workstation`.

```
[student@workstation ~]$ curl http://servera.lab.example.com:82
Hello
```

- ▶ 8. Return to the `workstation` system as the `student` user.

```
[root@servera ~]# exit
logout
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish netsecurity-ports
```

► Lab

Manage Network Security

Configure firewall and SELinux settings to allow access to multiple web servers that run on the same host.

Outcomes

- Configure firewall and SELinux settings on a web server host.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start netsecurity-review
```

Instructions

Your company decided to run a new web application. This application listens on the `80/TCP` and `1001/TCP` ports. All changes that you make must persist across a reboot.



Important

The Red Hat Online Learning environment needs the `5900/TCP` port to remain available to use the graphical interface. This port is also known under the `vnc-server` service. If you accidentally lock yourself out from the `serverb` machine, then you can either try to recover by using the `ssh` command to your `serverb` machine from your `workstation` machine, or reset your `serverb` machine. If you elect to reset your `serverb` machine, then you must run the setup scripts for this lab again. The configuration on your machines already includes a custom zone called `R0L` that opens these ports.

1. From the `workstation` machine, test access to the default web server at `http://serverb.lab.example.com` and to the `http://serverb.lab.example.com:1001` virtual host.
2. Log in to the `serverb` machine to determine what is preventing access to the web servers.
3. Configure SELinux to allow the `httpd` service to listen on the `1001/TCP` port.
4. From `workstation`, test again access to the default web server at `http://serverb.lab.example.com` and to the `http://serverb.lab.example.com:1001` virtual host.
5. Log in to the `serverb` machine to determine whether the correct ports are assigned to the firewall.

6. Add the 1001/TCP port to the permanent configuration for the `public` network zone. Confirm your configuration.
7. From `workstation`, confirm that the default web server at `http://serverb.lab.example.com` returns `SERVER B`, and that the virtual host at `http://serverb.lab.example.com:1001` returns `VHOST 1`.

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade netsecurity-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish netsecurity-review
```

► Solution

Manage Network Security

Configure firewall and SELinux settings to allow access to multiple web servers that run on the same host.

Outcomes

- Configure firewall and SELinux settings on a web server host.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start netsecurity-review
```

Instructions

Your company decided to run a new web application. This application listens on the `80/TCP` and `1001/TCP` ports. All changes that you make must persist across a reboot.



Important

The Red Hat Online Learning environment needs the `5900/TCP` port to remain available to use the graphical interface. This port is also known under the `vnc-server` service. If you accidentally lock yourself out from the `serverb` machine, then you can either try to recover by using the `ssh` command to your `serverb` machine from your `workstation` machine, or reset your `serverb` machine. If you elect to reset your `serverb` machine, then you must run the setup scripts for this lab again. The configuration on your machines already includes a custom zone called `R0L` that opens these ports.

1. From the `workstation` machine, test access to the default web server at `http://serverb.lab.example.com` and to the `http://serverb.lab.example.com:1001` virtual host.

- 1.1. Test access to the `http://serverb.lab.example.com` web server. The test currently fails. The web server should return `SERVER B`.

```
[student@workstation ~]$ curl http://serverb.lab.example.com
curl: (7) Failed to connect to serverb.lab.example.com port 80: Connection refused
```

- 1.2. Test access to the `http://serverb.lab.example.com:1001` virtual host. The test currently fails. The virtual host should return `VHOST 1`.

```
[student@workstation ~]$ curl http://serverb.lab.example.com:1001
curl: (7) Failed to connect to serverb.lab.example.com port 1001: No route to host
```

2. Log in to the `serverb` machine to determine what is preventing access to the web servers.

- 2.1. Log in to the `serverb` machine as the `student` user.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
[student@serverb ~]$
```

- 2.2. Determine whether the `httpd` service is active.

```
[student@serverb ~]$ systemctl is-active httpd
inactive
```

- 2.3. Enable and start the `httpd` service. The `httpd` service fails to start.

```
[student@serverb ~]$ sudo systemctl enable --now httpd
[sudo] password for student: student
Created symlink /etc/systemd/system/multi-user.target.wants/httpd.service → /usr/lib/systemd/system/httpd.service.
Job for httpd.service failed because the control process exited with error code.
See "systemctl status httpd.service" and "journalctl -xeu httpd.service" for details.
```

- 2.4. Investigate why the `httpd` service fails to start.

```
[student@serverb ~]$ systemctl status httpd.service
× httpd.service - The Apache HTTP Server
   Loaded: loaded (/usr/lib/systemd/system/httpd.service; enabled; vendor preset: disabled)
   Active: failed (Result: exit-code) since Wed 2022-04-13 06:55:01 EDT; 2min 52s ago
     Docs: man:httpd.service(8)
   Process: 1640 ExecStart=/usr/sbin/httpd $OPTIONS -DFOREGROUND (code=exited, status=1/FAILURE)
   Main PID: 1640 (code=exited, status=1/FAILURE)
   Status: "Reading configuration..."
     CPU: 31ms

Apr 13 06:55:01 serverb.lab.example.com systemd[1]: Starting The Apache HTTP Server...
Apr 13 06:55:01 serverb.lab.example.com httpd[1640]: (13)Permission denied: AH00072: make_sock: could not bind to address [::]:1001
Apr 13 06:55:01 serverb.lab.example.com httpd[1640]: (13)Permission denied: AH00072: make_sock: could not bind to address 0.0.0.0:1001
Apr 13 06:55:01 serverb.lab.example.com httpd[1640]: no listening sockets available, shutting down
Apr 13 06:55:01 serverb.lab.example.com httpd[1640]: AH00015: Unable to open logs
```

```
Apr 13 06:55:01 serverb.lab.example.com systemd[1]: httpd.service: Main process
exited, code=exited, status=1/FAILURE
Apr 13 06:55:01 serverb.lab.example.com systemd[1]: httpd.service: Failed with
result 'exit-code'.
Apr 13 06:55:01 serverb.lab.example.com systemd[1]: Failed to start The Apache
HTTP Server.
```

- 2.5. Check whether SELinux is blocking the httpd service from binding to the 1001/TCP port.

```
[student@serverb ~]$ sudo sealert -a /var/log/audit/audit.log
100% done
found 1 alerts in /var/log/audit/audit.log
-----

SELinux is preventing /usr/sbin/httpd from name_bind access on the tcp_socket port
1001.

**** Plugin bind_ports (99.5 confidence) suggests ****

If you want to allow /usr/sbin/httpd to bind to network port 1001
Then you need to modify the port type.
Do
# semanage port -a -t PORT_TYPE -p tcp 1001
   where PORT_TYPE is one of the following: http_cache_port_t, http_port_t,
   jboss_management_port_t, jboss_messaging_port_t, ntop_port_t, puppet_port_t.

**** Plugin catchall (1.49 confidence) suggests ****
...output omitted...
```

3. Configure SELinux to allow the httpd service to listen on the 1001/TCP port.

- 3.1. Use the semanage command to find the correct port type.

```
[student@serverb ~]$ sudo semanage port -l | grep 'http'
http_cache_port_t      tcp      8080, 8118, 8123, 10001-10010
http_cache_port_t      udp      3130
http_port_t          tcp      80, 81, 443, 488, 8008, 8009, 8443, 9000
pegasus_http_port_t    tcp      5988
pegasus_https_port_t   tcp      5989
```

- 3.2. Bind the 1001/TCP port to the http_port_t type.

```
[student@serverb ~]$ sudo semanage port -a -t http_port_t -p tcp 1001
```

- 3.3. Confirm that the 1001/TCP port is bound to the http_port_t port type.

```
[student@serverb ~]$ sudo semanage port -l | grep '^http_port_t'
http_port_t          tcp      1001, 80, 81, 443, 488, 8008, 8009, 8443, 9000
```

- 3.4. Enable and start the httpd service.


```
[student@serverb ~]$ sudo systemctl enable --now httpd
```

3.5. Verify the running state of the `httpd` service.

```
[student@serverb ~]$ systemctl is-active httpd
active
[student@serverb ~]$ systemctl is-enabled httpd
enabled
```

3.6. Return to the workstation machine as the `student` user.

```
[student@serverb ~]$ exit
logout
Connection to serverb closed.
[student@workstation ~]$
```

4. From workstation, test again access to the default web server at `http://serverb.lab.example.com` and to the `http://serverb.lab.example.com:1001` virtual host.

4.1. Test access to the `http://serverb.lab.example.com` web server. The web server should return `SERVER B`.

```
[student@workstation ~]$ curl http://serverb.lab.example.com
SERVER B
```

4.2. Test access to the `http://serverb.lab.example.com:1001` virtual host. The test continues to fail.

```
[student@workstation ~]$ curl http://serverb.lab.example.com:1001
curl: (7) Failed to connect to serverb.lab.example.com port 1001: No route to host
```

5. Log in to the `serverb` machine to determine whether the correct ports are assigned to the firewall.

5.1. Log in to the `serverb` machine as the `student` user.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
[student@serverb ~]$
```

5.2. Verify that the default firewall zone is set to the `public` zone.

```
[student@serverb ~]$ firewall-cmd --get-default-zone
public
```

5.3. If the previous step does not return `public` as the default zone, then correct it with the following command:

```
[student@serverb ~]$ sudo firewall-cmd --set-default-zone public
```

- 5.4. Determine the open ports that are listed in the `public` network zone.

```
[student@serverb ~]$ sudo firewall-cmd --zone=public --list-all
[sudo] password for student: student
public
  target: default
  icmp-block-inversion: no
  interfaces:
  sources:
  services: cockpit dhcpv6-client http ssh
  ports:
  protocols:
  forward: yes
  masquerade: no
  forward-ports:
  source-ports:
  icmp-blocks:
  rich rules:
```

6. Add the `1001/TCP` port to the permanent configuration for the `public` network zone. Confirm your configuration.

- 6.1. Add the `1001/TCP` port to the `public` network zone.

```
[student@serverb ~]$ sudo firewall-cmd --permanent --zone=public \
--add-port=1001/tcp
success
```

- 6.2. Reload the firewall configuration.

```
[student@serverb ~]$ sudo firewall-cmd --reload
success
```

- 6.3. Verify your configuration.

```
[student@serverb ~]$ sudo firewall-cmd --zone=public --list-all
public
  target: default
  icmp-block-inversion: no
  interfaces:
  sources:
  services: cockpit dhcpv6-client http ssh
  ports: 1001/tcp
  protocols:
  forward: yes
  masquerade: no
  forward-ports:
  source-ports:
  icmp-blocks:
  rich rules:
```

- 6.4. Return to the workstation machine as the `student` user.

```
[student@serverb ~]$ exit
logout
Connection to serverb closed.
[student@workstation ~]$
```

7. From workstation, confirm that the default web server at `http://serverb.lab.example.com` returns `SERVER B`, and that the virtual host at `http://serverb.lab.example.com:1001` returns `VHOST 1`.

7.1. Test access to the `http://serverb.lab.example.com` web server.

```
[student@workstation ~]$ curl http://serverb.lab.example.com
SERVER B
```

7.2. Test access to the `http://serverb.lab.example.com:1001` virtual host.

```
[student@workstation ~]$ curl http://serverb.lab.example.com:1001
VHOST 1
```

Evaluation

As the student user on the workstation machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade netsecurity-review
```

Finish

On the workstation machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish netsecurity-review
```

Summary

- The `netfilter` framework enables kernel modules to inspect every packet that traverses the system, including all incoming, outgoing, or forwarded network packets.
- The `firewalld` service simplifies management by classifying all network traffic into zones. Each zone has its own list of ports and services. The `public` zone is set as the default zone.
- The `firewalld` service ships with predefined services. You can list these services by using the `firewall-cmd --get-services` command.
- SELinux policy controls network traffic by labeling the network ports. For example, the `ssh_port_t` label is associated with the 22/TCP port. When a process wants to listen on a port, SELinux verifies whether the port's associated label is allowed to bind that port label.
- Use the `semanage` command to add, delete, and modify labels.

Chapter 12

Install Red Hat Enterprise Linux

Goal

Install Red Hat Enterprise Linux on servers and virtual machines.

Objectives

- Install Red Hat Enterprise Linux on a server.
- Explain Kickstart concepts and architecture, create a Kickstart file with the Kickstart Generator website, modify an existing Kickstart file with a text editor and check its syntax with `ksvalidator`, publish a Kickstart file to the installer, and install Kickstart on the network.
- Install a virtual machine on your Red Hat Enterprise Linux server with the web console.

Sections

- Install Red Hat Enterprise Linux (and Guided Exercise)
- Automate Installation with Kickstart (and Guided Exercise)
- Install and Configure Virtual Machines (and Quiz)

Lab

- Install Red Hat Enterprise Linux

Install Red Hat Enterprise Linux

Objectives

- Install Red Hat Enterprise Linux on a server.

Installation Media

Red Hat provides different forms of installation media that you can download from the Customer Portal website by using your active subscription.

- A binary image file in ISO 9660 format that contains the *Anaconda* Red Hat Enterprise Linux installation program, and the BaseOS and AppStream package repositories. These repositories contain the needed packages to complete the installation without additional repositories.
- A smaller "boot ISO" image file that contains Anaconda requires a configured network to access package repositories that are made available by using HTTP, FTP, or NFS.
- A QCOW2 image contains a prebuilt system disk that is ready to deploy as a virtual machine in cloud or enterprise virtual environments. Red Hat uses QCOW2 as the standard image format for KVM-based virtualization.
- Source code (human-readable programming language instructions) for Red Hat Enterprise Linux. The source DVDs have no documentation. This image helps to compile or develop your software according to the Red Hat Enterprise Linux version.

Red Hat Enterprise Linux 9 supports the following architectures:

- AMD and Intel 64-bit architectures (x86-64-v2)
- The 64-bit ARM architecture (ARMv8.0-A)
- 64-bit IBM Z (z14)
- IBM Power Systems, Little Endian (POWER9)

After downloading, create bootable installation media according to the instructions in the reference section.

Build Images by Using Image Builder

The Red Hat Image Builder tool helps to create customized images of Red Hat Enterprise Linux. Image Builder enables administrators to build custom system images for deployment on cloud platforms or on virtual environments for specialized use cases.

Use the `composer -cli` command or the Red Hat web console interface to access Image Builder.

Disk and Memory Requirements

To install Red Hat Enterprise Linux 9, you must have a minimum of 10 GiB of space on the partition where RHEL will be installed.

The minimum RAM requirement depends on the installation type:

Installation type	Recommended minimum RAM size
Local media installation (USB, DVD)	1.5 GiB for x86-64-v2, ARMv8.0-A, and z14 architectures 3 GiB for POWER9 architecture
NFS network installation	1.5 GiB for x86-64-v2, ARMv8.0-A, and z14 architectures 3 GiB for POWER9 architecture
HTTP, HTTPS, or FTP network installation	3 GiB for x86-64-v2, and z14 architectures 4 GiB for ARMv8.0-A, and POWER9 architecture

Install Red Hat Enterprise Linux Manually

By using the binary DVD or boot ISO, administrators install a new RHEL system on a bare-metal server or on a virtual machine. The Anaconda program supports two installation methods: manual and automated.

- The manual installation interacts with the user to query how Anaconda installs and configures the system.
- The automated installation uses a *Kickstart* file to direct Anaconda how to install the system.

Install RHEL by Using the Graphical Interface

Anaconda starts as a graphical application when you boot the system from the binary DVD or from the boot ISO.

At the **WELCOME TO RED HAT ENTERPRISE LINUX 9** screen, select the language, and click **Continue**. Individual users can choose a preferred language after installation.

Anaconda presents the **INSTALLATION SUMMARY** window, the central interface to customize parameters before beginning the installation.




Figure 12.1: Installation summary window

From this window, configure the installation parameters by selecting the icons in any order. Select an item to view or to edit. In any item, click **Done** to return to this central screen.

Anaconda marks mandatory items with a triangle warning symbol and message. The orange status bar at the bottom of the screen reminds you to complete the required information before the installation begins.

Complete the following items as needed:

- **Keyboard:** Add keyboard layouts.
- **Language Support:** Select additional languages to install.
- **Time & Date:** Select the system's location city by clicking the interactive map or selecting it from the lists. Specify the local time zone even when using *Network Time Protocol (NTP)*.
- **Connect to Red Hat:** Register the system with your Red Hat account and select the *system purpose*. The system purpose feature enables the registration process to automatically attach the most appropriate subscription to the system. You must first connect to the network by using the **Network & Host Name** icon to register the system.
- **Installation Source:** Provide the source package location that Anaconda requires for installation. The installation source field already refers to the DVD when using the binary DVD.
- **Software Selection:** Select the base environment to install, and add any add-ons. The **Minimal Install** environment installs only the essential packages to run Red Hat Enterprise Linux.
- **Installation Destination:** Select and partition the disks for Red Hat Enterprise Linux to install to. To complete this task, the administrator must know partitioning schemes and file-system

- selection criteria. The default radio button for automatic partitioning allocates the selected storage devices by using all available space.
- **KDUMP:** The *kdump* kernel crash dump feature collects information about the state of the system memory when the kernel crashes. Red Hat engineers analyze a *kdump* file to identify the cause of a crash. Use this Anaconda item to enable or to disable *kdump*.
 - **Network & Host Name:** Detected network connections are listed on the left. Select a connection to display its details. By default, Anaconda activates the network automatically. Click **Configure** for the selected network connection.
 - **Security Profile:** By activating a security profile, Anaconda applies restrictions and recommendations that the selected profile defines during installation.
 - **Root Password:** The installation program prompts to set a **root** password. The final stage of the installation process continues only after you define a **root** password.
 - **User Creation:** Create an optional non-root account. Creating a local, general-use account is a recommended practice. You can also create accounts after the installation is complete.



Note

When setting the **root** user password, Red Hat Enterprise Linux 9 enables an option to lock the **root** user access to the system. Red Hat Enterprise Linux 9 also enables password-based SSH access to the **root** user.

After you complete the installation configuration, and resolve all warnings, click **Begin Installation**. Clicking **Quit** aborts the installation without applying any changes to the system.

When the installation finishes, click **Reboot**. Anaconda displays the **Initial Setup** screen when installing a graphical desktop. Accept the license information and optionally register the system with the subscription manager. You might skip system registration until later.

Troubleshoot the Installation

During a Red Hat Enterprise Linux 9 installation, Anaconda provides two virtual consoles. The first virtual console has five windows that the **tmux** software terminal multiplexer supplies. You can access that console by using **Ctrl+Alt+F1**. The second virtual console, which is displayed by default, shows the Anaconda graphical interface. You can access it by using **Ctrl+Alt+F6**.

The **tmux** terminal provides a shell prompt in the second window in the first virtual console. You can use the terminal to enter commands to inspect and troubleshoot the system while the installation continues. The other windows provide diagnostic messages, logs, and additional information.

The following table lists the keystroke combinations to access the virtual consoles and the **tmux** terminal windows. In the **tmux** terminal, the keyboard shortcuts are performed in two actions: press and release **Ctrl+B**, and then press the number key of the window to access. In the **tmux** terminal, you can also use **Alt+Tab** to rotate the current focus between the windows.

Key sequence	contents
Ctrl+Alt+F1	Access the tmux terminal multiplexer.
Ctrl+B 1	In the tmux terminal, access the main information page for the installation process.

Key sequence	contents
Ctrl+B 2	In the <code>tmux</code> terminal, provide a root shell. Anaconda stores the installation log files in the <code>/tmp</code> directory.
Ctrl+B 3	In the <code>tmux</code> terminal, display the contents of the <code>/tmp/anaconda.log</code> file.
Ctrl+B 4	In the <code>tmux</code> terminal, display the contents of the <code>/tmp/storage.log</code> file.
Ctrl+B 5	In the <code>tmux</code> terminal, display the contents of the <code>/tmp/program.log</code> file.
Ctrl+Alt+F6	Access the Anaconda graphical interface.

**Note**

For compatibility with earlier Red Hat Enterprise Linux versions, the virtual consoles from Ctrl+Alt+F2 through Ctrl+Alt+F5 also present root shells during installation.

**References**

For further information, refer to *Understanding the Various RHEL .iso Files* at <https://access.redhat.com/solutions/104063>

For further information, refer to *Creating a Bootable Installation Medium for Red Hat Enterprise Linux* at

https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/performing_a_standard_rhel_9_installation/index#assembly_creating-a-bootable-installation-medium_assembly_preparing-for-your-installation

For further information, refer to *Composing a Customized RHEL System Image* at

https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/composing_a_customized_rhel_system_image/index

For further information, refer to *Performing a Standard RHEL 9 Installation* at

https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/performing_a_standard_rhel_9_installation/index

► Guided Exercise

Install Red Hat Enterprise Linux

Reinstall one of your servers with a minimal installation of Red Hat Enterprise Linux.

Outcomes

- Manually install Red Hat Enterprise Linux 9.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start installing-install
```

Instructions

- 1. Access the `servera` console and reboot the system into the installation media.
 - 1.1. Locate the `servera` console icon in your classroom environment. Open the console.
 - 1.2. To reboot, send `Ctrl+Alt+Del` to your system by using the relevant keyboard, virtual, or menu entry.
 - 1.3. When the boot loader menu appears, select the **Install Red Hat Enterprise Linux 9** menu entry.
- 2. Keep the default selected language and click **Continue**.
- 3. Click **Network & Host Name** to configure the network.
 - 3.1. Enter `servera.lab.example.com` in the **Host Name** field and then click **Apply**.
 - 3.2. Click **Configure** and then click the **IPv4 Settings** tab.
 - 3.3. Enter the following details:

Field	Value
Address	172.25.250.10
Netmask	24
Gateway	172.25.250.254
DNS servers	172.25.250.254

**Note**

Because Red Hat Enterprise Linux is already installed on the `servera.lab.example.com` machine, the values for each of the required fields exist. For a clean server, you must complete this information.

- 3.4. Click **Save** to save the network configuration, and then click **Done**.
- ▶ 4. Click **Installation Source** to configure the network installation source.
 - 4.1. Select the **On the network** option.
 - 4.2. Enter `content.example.com/rhel9.3/x86_64/dvd` in the network field.
 - 4.3. Click **Done**.
- ▶ 5. Click **Installation Destination** to select the installation disk.
 - 5.1. Select the `vda` disk and then click **Done**.
 - 5.2. On the next page, click **Reclaim space**.

**Note**

The `/dev/vda` disk already has partitions and file systems from the previous installation. With this selection, you can wipe the disk for the new installation.

- 5.3. Click **Delete all**, and then click **Reclaim space**.
- ▶ 6. Click **Software Selection**, select **Minimal Install** from the **Base Environment** list, and then click **Done**.
- ▶ 7. Set `redhat` as the password for the `root` user.
 - 7.1. Click **Root Password** and enter `redhat` in the **Root Password** field.
 - 7.2. Enter `redhat` in the **Confirm** field.
 - 7.3. Click **Done** twice because the password fails the dictionary check.
- ▶ 8. Add the `student` user and set `student` as the password.
 - 8.1. Click **User Creation**.
 - 8.2. Enter `student` in the **Full Name** field. The **User name** field automatically fills `student` as the username.
 - 8.3. Select **Make this user administrator** to enable the `student` user to use the `sudo` command to run commands as the `root` user.
 - 8.4. Enter `student` in the **Password** field.
 - 8.5. Enter `student` in the **Confirm password** field.
 - 8.6. Click the **Done** button twice because the password fails the dictionary check.

- ▶ 9. Click **Begin Installation**.
- ▶ 10. Click **Reboot System** when the installation is complete.
- ▶ 11. When the system displays the login prompt, log in as the **student** user.
- ▶ 12. After you validate the installation, reset the **servera** machine from the web page of the classroom environment.

Finish

On the **workstation** machine, change to the **student** user home directory and use the **lab** command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish installing-install
```

Automate Installation with Kickstart

Objectives

- Explain Kickstart concepts and architecture, create a Kickstart file with the Kickstart Generator website, modify an existing Kickstart file with a text editor and check its syntax with `kvalidator`, publish a Kickstart file to the installer, and install Kickstart on the network.

Introduction to Kickstart

The *Kickstart* feature of Red Hat Enterprise Linux automates system installations. You can use Kickstart text files to configure disk partitioning, network interfaces, package selection, and customize the installation. The Anaconda installer uses Kickstart files for a complete installation without user interaction. The Kickstart feature is similar and uses an unattended installation answer file for Microsoft Windows.

Kickstart files begin with a list of commands that define how to install the target machine. The installer ignores comment lines, which start with the number sign (#) character. Additional sections each begin with a specific directive that starts with the percentage sign (%) character. The `%end` directive marks the end of each section.

The `%packages` section specifies which software to include on installation. Specify individual packages by name, without versions. The at sign (@) character denotes package groups (either by group or ID), and the `@^` characters denote environment groups (groups of package groups). Lastly, use the `@module:stream/profile` syntax to denote module streams.

Groups have mandatory, default, and optional components. Normally, Kickstart installs mandatory and default components. To exclude a package or a package group from the installation, precede it with a hyphen (-) character. Excluded packages or package groups might still install if they are mandatory dependencies of other requested packages.

A Kickstart configuration file typically includes one or more `%pre` and `%post` sections, which contain scripts that further configure the system. The `%pre` scripts execute before any disk partitioning is done. Typically, you use `%pre` scripts to initialize a storage or network device that the remainder of the installation requires. The `%post` scripts execute after the initial installation is complete. Scripts within the `%pre` and `%post` sections can use any available interpreter on the system, including Bash or Python. Avoid the use of a `%pre` section, because any errors that occur within it might be difficult to diagnose.

Lastly, you can specify as many sections as you need, in any order. For example, you can have two `%post` sections, and they are interpreted in order of appearance.



Note

The RHEL Image Builder is an alternative installation method to Kickstart files. Rather than a text file that provides installation instructions, Image Builder creates an image with all the required system changes. The RHEL Image Builder can create images for public clouds such as Amazon Web Services and Microsoft Azure, or for private clouds such as OpenStack or VMware. See this section's references for more information about the RHEL Image Builder.

Installation Commands

The following Kickstart commands configure the installation source and method:

- **url**: Specifies the URL that points to the installation media.

```
url --url="http://classroom.example.com/content/rhel9.0/x86_64/dvd/"
```

- **repo**: Specifies where to find additional packages for installation. This option must point to a valid DNF repository.

```
repo --name="appstream" --baseurl=http://classroom.example.com/content/rhel9.0/x86_64/dvd/AppStream/
```

- **text**: Forces a text mode installation.
- **vnc**: Enables the VNC viewer so you can access the graphical installation remotely over VNC.

```
vnc --password=redhat
```

Device and Partition Commands

The following Kickstart commands configure devices and partitioning schemes:

- **clearpart**: Removes partitions from the system before creating partitions.

```
clearpart --all --drives=vda,vdb
```

- **part**: Specifies the size, format, and name of a partition. Required unless the **autopart** or **mount** commands are present.

```
part /home --fstype=ext4 --label=homes --size=4096 --maxsize=8192 --grow
```

- **autopart**: Automatically creates a root partition, a swap partition, and an appropriate boot partition for the architecture. On large enough drives (50 GB+), this command also creates a /home partition.
- **ignoredisk**: Prevents Anaconda from modifying disks, and is useful alongside the **autopart** command.

```
ignoredisk --drives=sdc
```

- **boot loader**: Defines where to install the bootloader. Required.

```
bootloader --location=mbr --boot-drive=sda
```

- **volgroup**, **logvol**: Creates LVM volume groups and logical volumes.

```
part pv.01 --size=8192
volgroup myvg pv.01
logvol / --vgname=myvg --fstype=xfs --size=2048 --name=rootvol --grow
logvol /var --vgname=myvg --fstype=xfs --size=4096 --name=varvol
```


- **zerombr**: Initialize disks whose formatting is unrecognized.

Network Commands

The following Kickstart commands configure networking-related features:

- **network**: Configures network information for the target system. Activates network devices in the installer environment.

```
network --device=eth0 --bootproto=dhcp
```

- **firewall**: Defines the firewall configuration for the target system.

```
firewall --enabled --service=ssh,http
```

Location and Security Commands

The following Kickstart commands configure security, language, and region settings:

- **lang**: Sets the language to use on the installed system. Required.

```
lang en_US
```

- **keyboard**: Sets the system keyboard type. Required.

```
keyboard --vckeymap=us
```

- **timezone**: Defines the time zone and whether the hardware clock uses UTC. Required.

```
timezone --utc Europe/Amsterdam
```

- **timesource**: Enables or disables NTP. If you enable NTP, then you must specify NTP servers or pools.

```
timesource --ntp-server classroom.example.com
```

- **authselect**: Sets up authentication options. Options that the **authselect** command recognizes are valid for this command. See **authselect(8)**.
- **rootpw**: Defines the initial **root** user password. Required.

```
rootpw --plaintext redhat
or
rootpw --iscrypted $6$KUnFfrTz08jv.PiH$YlBb0tXBkWzoMuRfb0.SpbQ...XDR1UuchoMG1
```

- **selinux**: Sets the SELinux mode for the installed system.

```
selinux --enforcing
```

- **services**: Modifies the default set of services to run under the default **systemd** target.

```
services --disabled=network, iptables, ip6tables --enabled=NetworkManager, firewalld
```

- **group, user:** Creates a local group or user on the system.

```
group --name=admins --gid=10001
user --name=jdoe --gecos="John Doe" --groups=admins
```

Miscellaneous Commands

The following Kickstart commands configure logging the host power state on completion:

- **logging:** This command defines how Anaconda handles logging during the installation.

```
logging --host=loghost.example.com
```

- **firstboot:** If enabled, then the Set up Agent starts the first time that the system boots. This command requires the `initial-setup` package.

```
firstboot --disabled
```

- **reboot, poweroff, halt:** Specify the final action when the installation completes. The default setting is the `halt` option.



Note

Most Kickstart commands have multiple available options. Review the *Kickstart Commands and Options* guide in this section's references for more information.

Example Kickstart File

In the following example, the first part of the Kickstart file consists of the installation commands, such as disk partitioning and installation source.

```
#version=RHEL9

# Define system bootloader options
bootloader --append="console=ttyS0 console=ttyS0,115200n8 no_timer_check
net.ifnames=0 crashkernel=auto" --location=mbr --timeout=1 --boot-drive=vda

# Clear and partition disks
clearpart --all --initlabel
ignoredisk --only-use=vda
zerombr
part / --fstype="xfs" --ondisk=vda --size=10000

# Define installation options
text
repo --name="appstream" --baseurl="http://classroom.example.com/content/rhel9.0/
x86_64/dvd/AppStream/"
url --url="http://classroom.example.com/content/rhel9.0/x86_64/dvd/"
```

```
# Configure keyboard and language settings
keyboard --vckeymap=us
lang en_US

# Set a root password, authselect profile, and selinux policy
rootpw --plaintext redhat
authselect select sssd
selinux --enforcing
firstboot --disable

# Enable and disable system services
services --disabled="kdump,rhsmcertd" --enabled="sshd,rngd,chronyd"

# Configure the system timezone and NTP server
timezone America/New_York --utc
timesource --ntp-server classroom.example.com
```

The second part of a Kickstart file contains the `%packages` section, with details of which packages and package groups to install, and which packages not to install.

```
%packages

@core
chrony
cloud-init
dracut-config-generic
dracut-norescue
firewalld
grub2
kernel
rsync
tar
-plymouth

%end
```

The last part of the Kickstart file contains a `%post` installation script.

```
%post

echo "This system was deployed using Kickstart on $(date)" > /etc/motd

%end
```

You can also specify a Python script with the `--interpreter` option.

```
%post --interpreter="/usr/libexec/platform-python"

print("This line of text is printed with python")

%end
```

**Note**

In a Kickstart file, missing required values cause the installer to interactively prompt for an answer or to abort the installation entirely.

Kickstart Installation Steps

Use the following steps to automate the installation of Red Hat Enterprise Linux with the Kickstart feature:

1. Create a Kickstart file.
2. Publish the Kickstart file so that the Anaconda installer can access it.
3. Boot the Anaconda installer and point it to the Kickstart file.

Create a Kickstart File

You can use either of these methods to create a Kickstart file:

- Use the Kickstart Generator website.
- Use a text editor.

The Kickstart Generator site at <https://access.redhat.com/labs/kickstartconfig> presents dialog boxes for user inputs, and creates a Kickstart configuration file with the user's choices. Each dialog box corresponds to the configurable items in the Anaconda installer.

The screenshot displays the Red Hat Customer Portal's Kickstart Generator. At the top, the Red Hat logo and 'Customer Portal' are visible, along with navigation links for Products & Services, Tools, Security, and Community. A search bar and language selector (English) are also present. The main heading is 'Kickstart Generator', with a subtext 'Generate a custom kickstart file based on your configuration parameters.' Below this, there's a link to 'Learn more about Kickstart Installations and how to use the generated config file in the Red Hat Enterprise Linux Install guide.' A dropdown menu shows 'Red Hat Enterprise Linux 9' with a 'DOWNLOAD' button. The 'Basic Configuration' section includes a sidebar with a list of configuration categories: Installation, Partition, BootLoader, Packages, Authentication, Network, Security, Display, Pre-Installation Script, and Post-Installation Script. The main configuration area contains several fields: 'Default Language' (English (United States)), 'Keyboard' (U.S. English), 'Time Zone' (America/New York), a checked checkbox for 'Use UTC clock', and two input fields for 'Root Password (required)' and 'Repeat Root Password (required)'.

Figure 12.2: The Red Hat Customer Portal Kickstart Generator

Creating a Kickstart file from scratch is complex, so first try to edit an existing file. Every installation creates a `/root/anaconda-ks.cfg` file that contains the Kickstart directives that are used in the installation. This file is a good starting point to create a Kickstart file.

The `ksvalidator` utility checks for syntax errors in a Kickstart file. It ensures that keywords and options are correctly used, but it does not validate URL paths, individual packages, groups, nor any part of `%post` or `%pre` scripts. For example, if the `firewall --disabled` directive is misspelled, then the `ksvalidator` command might produce one of the following errors:

```
[user@host ~]$ ksvalidator /tmp/anaconda-ks.cfg
The following problem occurred on line 10 of the kickstart file:

Unknown command: firewall

[user@host ~]$ ksvalidator /tmp/anaconda-ks.cfg
The following problem occurred on line 10 of the kickstart file:

no such option: --dsabled
```

The `ksverdiff` utility displays syntax differences between different operating system versions. For example, the following command displays the Kickstart syntax changes between RHEL 8 and RHEL 9:

```
[user@host ~]$ ksverdiff -f RHEL8 -t RHEL9
The following commands were removed in RHEL9:
device deviceprobe dmraid install multipath

The following commands were deprecated in RHEL9:
autostep btrfs method

The following commands were added in RHEL9:
timesource
...output omitted...
```

The `pykickstart` package provides the `ksvalidator` and `ksverdiff` utilities.

Publish the Kickstart File to Anaconda

Provide the Kickstart file to the installer by placing it in one of these locations:

- A network server that is available at installation time by using FTP, HTTP, or NFS.
- An available USB disk or CD-ROM.
- A local hard disk on the system.

The installer must access the Kickstart file to begin an automated installation. Usually, the file is made available via an FTP, web, or NFS server. Network servers help with Kickstart file maintenance, because changes can be made once, and then immediately be used for future installations.

By providing Kickstart files on USB or CD-ROM, is also convenient. The Kickstart file can be embedded in the boot media that starts the installation. However, when the Kickstart file is changed, you must generate new installation media.

Providing the Kickstart file on a local disk enables a quick rebuild of a system.

Boot Anaconda and Point to the Kickstart File

After a Kickstart method is chosen, the installer is told where to locate the Kickstart file by passing the `inst.ks=LOCATION` parameter to the installation kernel.

Consider the following examples:

- `inst.ks=http://server/dir/file`
- `inst.ks=ftp://server/dir/file`
- `inst.ks=nfs:server:/dir/file`
- `inst.ks=hd:device:/dir/file`
- `inst.ks=cdrom:device`




Figure 12.3: Specifying the Kickstart file location during installation

For virtual machine installations by using the **Virtual Machine Manager** or `virt-manager`, the Kickstart URL can be specified in a box under **URL Options**. When installing physical machines, boot by using installation media, and press the Tab key to interrupt the boot process. Add an `inst.ks=LOCATION` parameter to the installation kernel.



References

For further information, refer to *Kickstart Installation Basics* chapter in *Performing an Advanced RHEL Installation* at

https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/performing_an_advanced_rhel_9_installation/index#performing_an_automated_installation_using_kickstart

For further information, refer to *Kickstart Commands and Options Reference* in *Performing an Advanced RHEL Installation* at

https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/performing_an_advanced_rhel_9_installation/index#kickstart-commands-and-options-reference_installing-rhel-as-an-experienced-user

For further information, refer to *Composing a Customized RHEL System Image* at

https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/composing_a_customized_rhel_system_image/index

► Guided Exercise

Automate Installation with Kickstart

Create a kickstart file and validate the syntax.

Outcomes

- Create a kickstart file.
- Validate the kickstart file's syntax.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start installing-kickstart
```

Instructions

- 1. Log in to `servera` as the `student` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
```

- 2. Create the `/home/student/kickstart.cfg` file by copying the contents of the `/root/anaconda-ks.cfg` file using privileged access.

```
[student@servera ~]$ sudo cat /root/anaconda-ks.cfg > ~/kickstart.cfg
[sudo] password for student: student
```

- 3. Make the following changes to the `/home/student/kickstart.cfg` file.

- 3.1. Comment out the `reboot` command:

```
#reboot
```

- 3.2. Modify the `repo` commands to specify the content server's `BaseOS` and `AppStream` repositories:

```
repo --name="BaseOS" --baseurl="http://content.example.com/rhel9.3/x86_64/dvd/
BaseOS/"
repo --name="Appstream" --baseurl="http://content.example.com/rhel9.3/x86_64/dvd/
AppStream/"
```


- 3.3. Modify the `url` command to specify the content server's HTTP installation source:

```
url --url="http://content.example.com/rhel9.3/x86_64/dvd/"
```

- 3.4. Comment out the `network` command:

```
#network --bootproto=dhcp --device=link --activate
```

- 3.5. Modify the `rootpw` command to set the `root` user's password to `redhat`.

```
rootpw --plaintext redhat
```

- 3.6. Modify the `authselect` command to set the `sssd` service as the identity and authentication source.

```
authselect select sssd
```

- 3.7. Modify the `services` command to disable and enable services.

```
services --disabled="kdump,rhsmcertd" --enabled="sshd,rngd,chronyd"
```

- 3.8. Comment out the `part` commands and add the `autopart` command:

```
# Disk partitioning information
ignoredisk --only-use=vda
#part biosboot --fstype="biosboot" --size=1
#part /boot/efi --fstype="efi" --size=100 --fsoptions="..."
#part / --fstype="xfs"
autopart
```

- 3.9. Delete all of the content between the `%post` section and its `%end` directive. Add the `echo "Kickstarted on $(date)" >> /etc/issue` line.

```
%post --erroronfail
echo "Kickstarted on $(date)" >> /etc/issue
%end
```

- 3.10. Modify the `%packages` section to include only the following content:

```
%packages
@core
chrony
dracut-config-generic
dracut-norescue
firewalld
grub2
kernel
rsync
tar
```

```
httpd
-plymouth
%end
```

3.11. Save and exit the file.

- ▶ 4. Validate the Kickstart file for syntax errors. If no errors are shown, then the command has no output.

```
[student@servera ~]$ ksvalidator kickstart.cfg
```

- ▶ 5. Copy the `kickstart.cfg` file to the `/var/www/html/ks-config` directory.

```
[student@servera ~]$ sudo cp ~/kickstart.cfg /var/www/html/ks-config
[sudo] password for student: student
```

- ▶ 6. Return to the workstation machine as the `student` user.

```
[student@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the workstation machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish installing-kickstart
```

Install and Configure Virtual Machines

Objectives

- Install a virtual machine on your Red Hat Enterprise Linux server with the web console.

Introducing KVM Virtualization

Virtualization is a feature to support dividing a single physical machine into multiple *virtual machines* (VM), each of which can run an independent operating system.

Red Hat Enterprise Linux supports *KVM* (Kernel-based Virtual Machine), a full virtualization solution that is built into the standard Linux kernel. KVM can run multiple Windows and Linux guest operating systems.




Figure 12.4: KVM virtualization

In Red Hat Enterprise Linux, you can manage KVM from the command line with the `virsh` command, or graphically with the web console's Virtual Machines tool.

KVM virtual machine technology is available across all Red Hat products, from stand-alone physical instances of Red Hat Enterprise Linux to the Red Hat OpenStack Platform:

- Physical hardware systems run Red Hat Enterprise Linux to provide KVM virtualization. Red Hat Enterprise Linux is typically a *thick* host, a system that supports VMs and also providing other local and network services, applications, and management functions.
- *Red Hat Virtualization (RHV)* provides a centralized web interface that administrators can use to manage an entire virtual infrastructure. It includes advanced features such as KVM migration, redundancy, and high availability. A *Red Hat Virtualization Hypervisor* is a tuned version of Red Hat Enterprise Linux solely for provisioning and supporting VMs.
- *Red Hat OpenStack Platform (RHOSP)* provides the foundation to create, deploy, and scale a public or a private cloud.
- *Red Hat OpenShift Virtualization* includes RHV components to enable deployment of containers on bare metal.

On systems where SELinux is enabled, sVirt isolates guests and the hypervisor. Each virtual machine process is labeled and is automatically allocated a unique level, and the associated virtual disk files are given matching labels.

Virtual Machine Types

Two main virtual machine types are available to install the guest operating system on. The earlier of the two uses the i440FX chip set, and provides ISA and PCI bridges, an IDE controller, and VGA and Ethernet adapters. The Q35 machine type is the later of the two, and should be used in most cases. It supports later virtual hardware interfaces such as PCIe and SATA, and includes an SM Bus controller.

In Red Hat Enterprise Linux 8 and later versions, UEFI and Secure Boot support for virtual machines is provided by Open Virtual Machine Firmware (OVMF), in the `edk2-ovmf` package.

Supported Guest Operating Systems

For virtualization solutions, Red Hat provides guest operating systems that are based on the product that you are using.

Red Hat offers KVM for using single-machine virtualization on RHEL. For more scalable and advanced features, Red Hat offers Red Hat OpenShift Virtualization and Red Hat OpenStack Platform.

For more information about supported guest operating systems for RHV, RHOSP, and OpenShift Virtualization, refer to Certified Guest Operating Systems in Red Hat OpenStack Platform, Red Hat Virtualization, OpenShift Virtualization and Red Hat Enterprise Linux with KVM [<https://access.redhat.com/articles/973163>]

Configure a Red Hat Enterprise Linux Physical System as a Virtualization Host

Administrators can configure a Red Hat Enterprise Linux system as a virtualization host, as appropriate for development, testing, training, or when needing to work in multiple operating systems at the same time.

Install the Virtualization Tools

Install the `Virtualization Host` DNF package group to prepare a system to become a virtualization host.

```
[root@host ~]# dnf group list | grep -i virt
Virtualization Host
[root@host ~]# dnf group info "Virtualization Host"
...output omitted...
Environment Group: Virtualization Host
Description: Minimal virtualization host.
Mandatory Groups:
  Base
  Core
  Standard
  Virtualization Hypervisor
  Virtualization Tools
Optional Groups:
  Debugging Tools
  Network File System Client
  Remote Management for Linux
  Virtualization Platform
```

```
[root@host ~]# dnf group install "Virtualization Host"
...output omitted...
```

Verify the System Requirements for Virtualization

KVM requires either an Intel processor with the Intel VT-x and Intel 64 extensions for x86-based systems, or an AMD processor with the AMD-V and the AMD64 extensions. To verify your hardware and the system requirements, use the `virt-host-validate` command.

```
[root@host ~]# virt-host-validate
QEMU: Checking for hardware virtualization           : PASS
QEMU: Checking if device /dev/kvm exists             : PASS
QEMU: Checking if device /dev/kvm is accessible      : PASS
QEMU: Checking if device /dev/vhost-net exists       : PASS
QEMU: Checking if device /dev/net/tun exists         : PASS
QEMU: Checking for cgroup 'cpu' controller support   : PASS
QEMU: Checking for cgroup 'cpuacct' controller support : PASS
QEMU: Checking for cgroup 'cpuset' controller support : PASS
QEMU: Checking for cgroup 'memory' controller support : PASS
QEMU: Checking for cgroup 'devices' controller support : PASS
QEMU: Checking for cgroup 'blkio' controller support : PASS
QEMU: Checking for device assignment IOMMU support   : PASS
QEMU: Checking for secure guest support             : PASS
```

The system must pass all the validation items to operate as a KVM host.

Install a Virtual Machine from the Command Line

Install the `virt-install` package, and then use the `virt-install` command to deploy a virtual machine from the command line. Many installation sources are supported, such as ISO, NFS, HTTP, and FTP. The following example uses the `virt-install` command's `--cdrom` option to install the guest from an ISO image.

```
[root@host ~]# virt-install --name demo --memory 4096 --vcpus 2 --disk size=40 \
--os-type linux --cdrom /root/rhel.iso
...output omitted...
```

Manage Virtual Machines with the Web Console

The `libvirt-client` package provides the `virsh` command, to manage virtual machines from the command line. Alternatively, you can create and manage virtual machines graphically by using the web console Virtual Machines plug-in.

Install the `cockpit-machines` package to add the **Virtual Machines** menu to the web console.

```
[root@host ~]# dnf install cockpit-machines
```

If the web console is not already running, then start and enable it.

```
[root@host ~]# systemctl enable --now cockpit.socket
```

In a web browser that is running on the local machine, navigate to `https://localhost/9090` and log in to the web console. Switch to administrative mode and enter your password, if necessary.




Figure 12.5: Managing virtual machines in the web console

To create a virtual machine with the web console, access the **Virtual Machines** menu. From there, click **Create VM** and enter the VM configuration in the **Create New Virtual Machine** window. If you are using the web console for the first time after installing the Virtual Machines plug-in, then you must reboot your system to start the `libvirt` virtualization.

Figure 12.6: Creating a virtual machine in the web console

- **Name** sets a *domain* name for the virtual machine configuration. This name is unrelated to the hostname that you give the virtual machine during installation.
- **Installation type** is the method for accessing the installation media. Choices include the local file system, or an HTTPS, FTP, or NFS URL, or PXE.
- **Installation source** provides the path to the installation source.

- **Operating system** defines the virtual machine's operating system. The virtualization layer presents hardware emulation to be compatible with the chosen operating system.
- **Storage** defines whether to create a storage volume or to use an existing one.
- **Size** is the disk size when creating a new volume. Associate additional disks with the VM after installation.
- **Memory** is the amount of RAM to provide to the new VM.
- **Immediately start VM** indicates whether the VM immediately starts after you click **Create**.

Click **Create** to create the VM, and click **Install** to start the operating system installation. The web console displays the VM console from which you can install the system.



References

For more information, refer to the *Configuring and Managing Virtualization* guide at https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html-single/configuring_and_managing_virtualization/index

What Is Virtualization?

<https://www.redhat.com/en/topics/virtualization/what-is-virtualization>

Certified Guest Operating Systems in Red Hat OpenStack Platform, Red Hat Virtualization, OpenShift Virtualization and Red Hat Enterprise Linux with KVM

<https://access.redhat.com/articles/973163>

► Quiz

Install and Configure Virtual Machines

Choose the correct answers to the following questions:

- 1. Which package enables you to select the OVMF firmware for a virtual machine?
 - a. open-firmware
 - b. edk2-ovmf
 - c. core-ovmf
 - d. ovmf
 - e. virt-open-firmware

- 2. Which two components are required to configure your system as a virtualization host, and to manage virtual machines with the web console? (Choose two.)
 - a. The Virtualization Host package group
 - b. The openstack package group
 - c. The cockpit-machines package
 - d. The Virtualization Platform package group
 - e. The kvm DNF module
 - f. The cockpit-virtualization package

- 3. Which command verifies that your system supports virtualization?
 - a. `grep kvm /proc/cpuinfo`
 - b. `virsh validate`
 - c. `virt-host-validate`
 - d. `rhv-validate`
 - e. `cockpit-validate`

- 4. Which two tools can you use to start and stop your virtual machines on a Red Hat Enterprise Linux system? (Choose two.)
 - a. `vmctl`
 - b. `libvirt`
 - c. `virsh`
 - d. `neutron`
 - e. the web console

► Solution

Install and Configure Virtual Machines

Choose the correct answers to the following questions:

- 1. Which package enables you to select the OVMF firmware for a virtual machine?
 - a. open-firmware
 - b. edk2-ovmf
 - c. core-ovmf
 - d. ovmf
 - e. virt-open-firmware

- 2. Which two components are required to configure your system as a virtualization host, and to manage virtual machines with the web console? (Choose two.)
 - a. The Virtualization Host package group
 - b. The openstack package group
 - c. The cockpit-machines package
 - d. The Virtualization Platform package group
 - e. The kvm DNF module
 - f. The cockpit-virtualization package

- 3. Which command verifies that your system supports virtualization?
 - a. `grep kvm /proc/cpuinfo`
 - b. `virsh validate`
 - c. `virt-host-validate`
 - d. `rhv-validate`
 - e. `cockpit-validate`

- 4. Which two tools can you use to start and stop your virtual machines on a Red Hat Enterprise Linux system? (Choose two.)
 - a. `vmctl`
 - b. `libvirt`
 - c. `virsh`
 - d. `neutron`
 - e. the web console

► Lab

Install Red Hat Enterprise Linux

Create a kickstart file and make it available to the installer.

Outcomes

- Create a kickstart file.
- Provide the kickstart file to the installer.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start installing-review
```

Instructions

Prepare a kickstart file on the `serverb` machine as specified, and provide it at the `http://serverb.lab.example.com/ks-config/kickstart.cfg` address.

1. On the `serverb` machine, copy the `/root/anaconda-ks.cfg` kickstart file to the `/home/student/kickstart.cfg` kickstart file to be editable for the `student` user.
2. Update the `/home/student/kickstart.cfg` kickstart file.
 - Comment out the `reboot` command.
 - Modify the `repo` command for the BaseOS and AppStream repositories. Modify the `repo` command for the BaseOS repository to use the `http://content.example.com/rhel9.3/x86_64/dvd/BaseOS/` address. Modify the `repo` command for the AppStream repository to use the `http://content.example.com/rhel9.3/x86_64/dvd/AppStream/` address.
 - Change the `url` command to use `http://content.example.com/rhel9.3/x86_64/dvd/` as the installation source.
 - Comment out the `network` command.
 - Change the `rootpw` command to set `redhat` as the `root` user password.
 - Modify the `authselect` command to set the `sssd` service as the identity and authentication source.
 - Modify the `services` command to disable the `kdump` and `rhsmcertd` services and to enable the `sshd`, `rngd`, and `chronyd` services.
 - Delete the `part` commands and add the `autopart` command.

- Simplify the `%post` section so that it runs only a script to append the text `Kickstarted on DATE` at the end of the `/etc/issue` file. Use the `date` command to insert the date with no additional options.
 - Simplify the `%package` section as follows: include the `@core`, `chrony`, `dracut-config-generic`, `dracut-norescue`, `firewalld`, `grub2`, `kernel`, `rsync`, `tar`, and `httpd` packages. Ensure that the `plymouth` package does not install.
3. Validate the syntax of the `kickstart.cfg` kickstart file.
 4. Provide the `/home/student/kickstart.cfg` file at the `http://serverb.lab.example.com/ks-config/kickstart.cfg` address.

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade installing-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish installing-review
```

► Solution

Install Red Hat Enterprise Linux

Create a kickstart file and make it available to the installer.

Outcomes

- Create a kickstart file.
- Provide the kickstart file to the installer.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start installing-review
```

Instructions

Prepare a kickstart file on the `serverb` machine as specified, and provide it at the `http://serverb.lab.example.com/ks-config/kickstart.cfg` address.

1. On the `serverb` machine, copy the `/root/anaconda-ks.cfg` kickstart file to the `/home/student/kickstart.cfg` kickstart file to be editable for the `student` user.

- 1.1. Log in to the `servera` machine as the `student` user.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
[student@serverb ~]$
```

- 1.2. On the `serverb` machine, copy the `/root/anaconda-ks.cfg` file to the `/home/student/kickstart.cfg` file.

```
[student@serverb ~]$ sudo cat /root/anaconda-ks.cfg > ~/kickstart.cfg
[sudo] password for student: student
```

2. Update the `/home/student/kickstart.cfg` kickstart file.
 - Comment out the `reboot` command.
 - Modify the `repo` command for the BaseOS and AppStream repositories. Modify the `repo` command for the BaseOS repository to use the `http://content.example.com/rhel9.3/x86_64/dvd/BaseOS/` address. Modify the `repo` command for the AppStream repository to use the `http://content.example.com/rhel9.3/x86_64/dvd/AppStream/` address.

- Change the `url` command to use `http://content.example.com/rhel9.3/x86_64/dvd/` as the installation source.
- Comment out the `network` command.
- Change the `rootpw` command to set `redhat` as the `root` user password.
- Modify the `authselect` command to set the `sssd` service as the identity and authentication source.
- Modify the `services` command to disable the `kdump` and `rhsmcertd` services and to enable the `sshd`, `rngd`, and `chronyd` services.
- Delete the `part` commands and add the `autopart` command.
- Simplify the `%post` section so that it runs only a script to append the text `Kickstarted on DATE` at the end of the `/etc/issue` file. Use the `date` command to insert the date with no additional options.
- Simplify the `%package` section as follows: include the `@core`, `chrony`, `dracut-config-generic`, `dracut-norescue`, `firewalld`, `grub2`, `kernel`, `rsync`, `tar`, and `httpd` packages. Ensure that the `plymouth` package does not install.

2.1. Comment out the `reboot` directive.

```
#reboot
```

- 2.2. Modify the `repo` command for the BaseOS and AppStream repositories. Modify the `repo` command for the BaseOS repository to use the `http://content.example.com/rhel9.3/x86_64/dvd/BaseOS/` address. Modify the `repo` command for the AppStream repository to use the `http://content.example.com/rhel9.3/x86_64/dvd/AppStream/` address.

```
repo --name="BaseOS" --baseurl="http://content.example.com/rhel9.3/x86_64/dvd/BaseOS/"
repo --name="Appstream" --baseurl="http://content.example.com/rhel9.3/x86_64/dvd/AppStream/"
```

- 2.3. Change the `url` command to specify the HTTP installation source media that the classroom machine provides.

```
url --url="http://content.example.com/rhel9.3/x86_64/dvd/"
```

- 2.4. Comment out the `network` command.

```
#network --bootproto=dhcp --device=link --activate
```

- 2.5. Modify the `rootpw` command to set `redhat` as the password for the `root` user.

```
rootpw --plaintext redhat
```

- 2.6. Modify the `authselect` command to set the `sssd` service as the identity and authentication source.

```
authselect select sssd
```

2.7. Simplify the `services` command to match the following string.

```
services --disabled="kdump,rhsmcertd" --enabled="sshd,rngd,chronyd"
```

2.8. Comment out the `part` commands and add the `autopart` command:

```
# Disk partitioning information
ignoredisk --only-use=vda
#part biosboot --fstype="biosboot" --size=1
#part /boot/efi --fstype="efi" --size=100 --
fsoptions="defaults,uid=0,gid=0,umask=077,shortname=winnt"
#part / --fstype="xfs" --size=10137 --label=root
autopart
```

2.9. Delete all content between the `%post` and `%end` sections. Add the `echo "Kickstarted on $(date)" >> /etc/issue` line.

```
%post --erroronfail
echo "Kickstarted on $(date)" >> /etc/issue
%end
```

2.10. Simplify the package specification to match the following content.

```
%packages
@core
chrony
dracut-config-generic
dracut-norescue
firewalld
grub2
kernel
rsync
tar
httpd
-plymouth
%end
```

3. Validate the syntax of the `kickstart.cfg` kickstart file.

3.1. Validate the Kickstart file for syntax errors.

```
[student@serverb ~]$ ksvalidator kickstart.cfg
```

4. Provide the `/home/student/kickstart.cfg` file at the `http://serverb.lab.example.com/ks-config/kickstart.cfg` address.

4.1. Copy the `kickstart.cfg` file to the `/var/www/html/ks-config/` directory.

```
[student@serverb ~]$ sudo cp kickstart.cfg /var/www/html/ks-config/  
[sudo] password for student: student
```

4.2. Return to the `workstation` machine as the `student` user.

```
[student@serverb ~]$ exit  
logout  
Connection to serverb closed.  
[student@workstation ~]$
```

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade installing-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish installing-review
```

Summary

- The RHEL 9 binary DVD includes Anaconda and all required repositories for installation.
- The RHEL 9 boot ISO includes the Anaconda installer, and can access repositories over the network during installation.
- The Kickstart system can perform unattended installations.
- You can create Kickstart files by using the Kickstart Generator website or by copying and editing `/root/anaconda-ks.cfg`.
- The `Virtualization Host` DNF package group provides the packages for a RHEL system to become a virtualization host.
- The `cockpit-machines` package adds the **Virtual Machines** menu to Cockpit.

Chapter 13

Run Containers

Goal

Obtain, run, and manage simple lightweight services as containers on a single Red Hat Enterprise Linux server.

Objectives

- Explain container concepts and the core technologies for building, storing, and running containers.
- Pull container images from a registry and run containers by using Podman.
- Authenticate to and explore the container images stored in remote container registries.
- Manage the lifecycle of a container from creation to deletion.

Sections

- Introduction to Containers (and Quiz)
- Creating Containers with Podman (and Guided Exercise)
- Container Image Registries (and Guided Exercise)
- Managing the Container Lifecycle (and Guided Exercise)

Lab

- Run Containers

Introduction to Containers

Objectives

- Explain container concepts and the core technologies for building, storing, and running containers.

Describing Containers

In computing, a *container* is an encapsulated process that includes the required runtime dependencies for the program to run. In a container, application-specific libraries are independent of the host operating system libraries. Libraries and functions that are not specific to the containerized application are provided by the operating system and kernel. The provided libraries and functions help to ensure that the container remains compact, and that it can quickly execute and stop as needed.

A container engine creates a union file system by merging container image layers. Because container image layers are immutable, a container engine adds a writable layer for runtime file modifications. Containers are *ephemeral* by default, which means that the container engine removes the writable layer when you remove the container.




Figure 13.1: Applications in containers versus on host operating system

Containers use Linux kernel features, such as namespaces, control groups (cgroups), SELinux, and secure computing mode (seccomp). For example, containers use cgroups for resource management, such as CPU time allocation and system memory. Namespaces in particular provide the functionality to isolate processes within containers from each other and from the host system. As such, the environment within a container is Linux-based, regardless of the host operating system. When using containers on non-Linux operating systems, these Linux-specific features are often virtualized by the container engine implementation.

Containerization originated from technologies such as the `chroot` utility, a method to partially or fully isolate an environment, and evolved to the *Open Container Initiative (OCI)*, which is a governance organization that defines standards for creating and running containers. Most container engines conform to the OCI specifications, so developers can confidently build their deployable target artifacts to run as OCI containers.

Containers use SELinux and secure computing mode to enforce security boundaries and restrict the features available in containers.



Note

For a deeper discussion of container architecture and security, refer to the A Layered Approach to Container and Kubernetes Security [<https://www.redhat.com/en/resources/container-security-openshift-cloud-devops-whitepaper>] white paper.

Images versus Instances

Containers can be split into two similar but distinct ideas: *container images* and *container instances*. A *container image* contains effectively immutable data that defines an application and its libraries. You can use container images to create *container instances*, which are executable versions of the image that include references to networking, disks, and other runtime necessities.

You can use a single container image multiple times to create many distinct container instances. You can also run these instances across multiple hosts. The application within a container is independent of the host environment.



Note

OCI container images are defined by the `image-spec` specification, whereas OCI container instances are defined by the `runtime-spec` specification.

Another way to think about *container images* versus *container instances* is that an instance relates to an image as an object relates to a class in object-oriented programming.

Comparing Containers to Virtual Machines

Containers generally serve a similar role to *virtual machines* (VMs), where an application resides in a self-contained environment with virtualized networking for communication. Although this use case initially seems to be the same, containers have a smaller footprint, and start and stop faster than a virtual machine. For both memory and disk usage, VMs are often measured in gigabytes, whereas containers are measured in megabytes.

A VM is useful when an additional full computing environment is required, such as when an application requires specific, dedicated hardware. Additionally, a VM is preferable when an application requires a non-Linux operating system or a different kernel from the host.

Virtual Machines versus Containers

Virtual machines and containers use different software for management and functionality. Hypervisors, such as KVM, Xen, VMware, and Hyper-V, are applications that provide the virtualization functionality for VMs. The container equivalent of a hypervisor is a container engine, such as Podman.

	Virtual machines	Containers
Machine-level functionality	Hypervisor	Container engine

	Virtual machines	Containers
Management	VM management interface	Container engine or orchestration software
Virtualization level	Fully virtualized environment	Only relevant parts
Size	Measured in gigabytes	Measured in megabytes
Portability	Generally only same hypervisor	Any OCI-compliant engine

You can manage hypervisors with additional management software. The software can be included with the hypervisor, or it can be external, such as Virtual Machine Manager with KVM. Alternatively, you can manage containers directly through the container engine itself. Additionally, you can use container orchestration tools, such as Red Hat OpenShift Container Platform (RHOCP) and Kubernetes, to run and manage containers at scale. RHOCP manages both containers and virtual machines from a common interface.

With VMs, interoperability is uncommon. A VM that runs on one hypervisor is usually not guaranteed to run on a different hypervisor. In contrast, containers that follow the OCI specification do not require a particular container engine to function. Many container engines can function as drop-in replacements for each other.

Deployment at Scale

Both containers and VMs can work well at various scales. Because a container requires considerably fewer resources than a VM, containers have performance and resource benefits at a larger scale. A common method in large-scale environments is to use containers that run inside VMs. This configuration takes advantage of the strong points in each technology.

Development for Containers

Containerization provides many advantages for the development process, such as easier testing and deployments, by providing tools for stability, security, and flexibility.

Testing and Workflows

One of the greatest advantages of containers for developers is the ability to scale. A developer can write software and test locally, and then deploy the finished application to a cloud server or a dedicated cluster with few or no changes. This workflow is especially useful when creating microservices, which are small and ephemeral containers that are designed to spin up and down as needed. Additionally, developers that use containers can take advantage of *Continuous Integration/Continuous Development (CI/CD)* pipelines to deploy containers to various environments. In particular, RHOCP offers various integration features with CI/CD pipelines and workflows in mind.

Stability

Container images are a stable target for developers. Software applications require specific versions of libraries to be available for deployment, which can result in dependency issues or specific OS requirements. Because libraries are included in the container image, a developer can be confident of no dependency issues in a deployment. Having the libraries integrated within the container removes variability between testing and production environments. For example,

a container with a specific version of Python ensures that the same version of Python is used in every testing or deployment environment.

Multi-container Applications

A multi-container application is distributed across several containers. You can run the containers from the same image for *high availability* (HA) replicas, or from several different images. For example, a developer can create an application that includes a database container that runs separately from the application's web API container. An application can rely on the container management software to provide HA replicas with multi-container options, such as Podman Pods or the `compose - spec` specification.



References

Red Hat Technology Topics: Understanding Containers

<https://www.redhat.com/en/topics/containers>

Red Hat Technology Topics: Understanding Virtualization

<https://www.redhat.com/en/topics/virtualization>

Podman Documentation

<https://docs.podman.io/en/latest/>

About the Open Container Initiative

<https://opencontainers.org/about/overview/>

Kubernetes Documentation

<https://kubernetes.io/docs/home/>

For more information about Red Hat OpenShift Container Platform, refer to the documentation at

https://access.redhat.com/documentation/en-us/openshift_container_platform/4.14

► Quiz

Introduction to Containers

Choose the correct answers to the following questions:

- 1. **What is an effectively immutable version of an application that includes its libraries?**
 - a. Container instance
 - b. Virtual machine
 - c. Container image
 - d. Container engine
- 2. **What are two advantages of using containers? (Choose two.)**
 - a. Faster program execution
 - b. Stable development target
 - c. Ease of testing and deployment
 - d. Fully virtualized environment
- 3. **Which organization directly governs the standards for containerization?**
 - a. Docker
 - b. Red Hat
 - c. Open Container Initiative
 - d. Linux Foundation
- 4. **True or false? An application in an OCI container runs on any OCI-compliant container engine.**
 - a. True
 - b. False
- 5. **What is an executable version of an application that includes references to networking, disks, and other runtime necessities?**
 - a. Container instance
 - b. Container image
 - c. Container engine
 - d. Virtual machine
- 6. **True or false? A virtual machine most often has a smaller footprint compared to a container.**
 - a. True
 - b. False

- 7. What provides machine-level functionality for running containerized applications?
- a. Container instance
 - b. Container image
 - c. Container engine
 - d. Hypervisor

► Solution

Introduction to Containers

Choose the correct answers to the following questions:

- 1. **What is an effectively immutable version of an application that includes its libraries?**
 - a. Container instance
 - b. Virtual machine
 - c. Container image
 - d. Container engine
- 2. **What are two advantages of using containers? (Choose two.)**
 - a. Faster program execution
 - b. Stable development target
 - c. Ease of testing and deployment
 - d. Fully virtualized environment
- 3. **Which organization directly governs the standards for containerization?**
 - a. Docker
 - b. Red Hat
 - c. Open Container Initiative
 - d. Linux Foundation
- 4. **True or false? An application in an OCI container runs on any OCI-compliant container engine.**
 - a. True
 - b. False
- 5. **What is an executable version of an application that includes references to networking, disks, and other runtime necessities?**
 - a. Container instance
 - b. Container image
 - c. Container engine
 - d. Virtual machine
- 6. **True or false? A virtual machine most often has a smaller footprint compared to a container.**
 - a. True
 - b. False

► 7. What provides machine-level functionality for running containerized applications?

- a. Container instance
- b. Container image
- c. Container engine
- d. Hypervisor

Creating Containers with Podman

Objectives

- Pull container images from a registry and run containers by using Podman.

An Introduction to Podman

Podman is an open source tool that you can use to manage your containers locally. With Podman, you can find, run, build, or deploy OCI (Open Container Initiative) containers and container images.

By default, Podman is daemonless. A daemon is a process that is always running and ready to receive incoming requests. Some other container tools use a daemon to proxy the requests, which brings a single point of failure. In addition, a daemon might require elevated privileges, which is a security concern. Podman interacts directly with containers, images, and registries without a daemon.

Podman comes in the form of a command-line interface (CLI), which is supported for several operating systems. Along with the CLI, Podman provides two additional ways to interact with your containers and to automate processes: the RESTful API and a desktop application called *Podman Desktop*.

Working with Podman

After you install Podman, you can use it by running the `podman` command. The following command displays the version that you are using.

```
[user@host ~]$ podman -v
podman version VERSION
```

Pulling and Displaying Images

Before you can run your application in a container, you must create a container image.

With Podman, you fetch container images from image registries by using the `podman pull` command. For example, the following command fetches a containerized version of Red Hat Enterprise Linux 9 from the Red Hat Registry.

```
[user@host ~]$ podman pull registry.redhat.io/rhel9/rhel-guest-image:9.4
Trying to pull registry.redhat.io/rhel9/rhel-guest-image:9.4...
Getting image source signatures
...output omitted...
Writing manifest to image destination
Storing signatures
b85986059f7663c1b89431f74cdcb783f6540823e4b85c334d271f2f2d8e06d6
```

The containerized RHEL 9 image is referenced in the form *NAME:VERSION*. In the following example, you fetch the 9.4 version of the `registry.redhat.io/rhel9/rhel-guest-image` image.

After you run the `pull` command, the image is stored locally in your system. You can list the images in your system by using the `podman images` command.

```
[user@host ~]$ podman images
REPOSITORY                                TAG          IMAGE ID      CREATED      SIZE
registry.redhat.io/rhel9/rhel-guest-image  9.4          52617ef413bd  2 weeks ago  216 MB
```

Container and Container Images

A container is an isolated runtime environment where applications are executed as isolated processes. The isolation of the runtime environment ensures that they do not interrupt other containers or system processes.

A container image contains a packaged version of your application, with all the dependencies that are necessary for the application to run. Images can exist without containers, but containers are dependent on images because containers use container images to build a runtime environment to execute applications.

Running and Displaying Containers

With the image stored in your local system, you can use the `podman run` command to create a new container that uses the image. The RHEL image from previous examples accepts Bash commands as an argument. These commands are provided as an argument and are executed within a RHEL container.

```
[user@host ~]$ podman run registry.redhat.io/rhel9/rhel-guest-image:9.4 \
echo 'Red Hat'
Red Hat
```

In the previous example, the `echo 'Red Hat'` command is provided as an argument to the `podman run` command. Podman executes the `echo` command inside the RHEL container and displays the output of the command.



Note If you run a container from an image that is not stored in your system, then Podman tries to pull the image from a container registry before running the container. Therefore, it is not necessary to execute the `pull` command first.

When the container finishes executing the `echo` command, the container stops because no other process keeps it running. You can list the running containers by using the `podman ps` command.

```
[user@host ~]$ podman ps
CONTAINER ID  IMAGE          COMMAND        CREATED        STATUS        PORTS        NAMES
```

By default, the `podman ps` command lists the following details for your containers.

- The container's ID

- The name of the image that the container is using
- The command that the container is executing
- The time that the container was created
- The status of the container
- The exposed ports in the container
- The name of the container

However, *stopping* a container is not the same as *removing* a container. Although the container is stopped, Podman does not remove it. You can list all containers (running and stopped) by adding the `--all` flag to the `podman ps` command.

```
[user@host ~]$ podman ps --all
CONTAINER ID  IMAGE                                COMMAND                                CREATED
STATUS        PORTS        NAMES
20236410bcef  registry.redhat.io/rhel9/rhel-guest-image:9.4  echo Red Hat  1
second ago    Exited (0) 1 second ago                hungry_mclaren
```

You can also automatically remove a container when it exits by adding the `--rm` option to the `podman run` command.

```
[user@host ~]$ podman run --rm registry.redhat.io/rhel9/rhel-guest-image:9.4 \
echo 'Red Hat'
Red Hat
[user@host ~]$ podman ps --all
CONTAINER ID  IMAGE                                COMMAND                                CREATED
STATUS        PORTS        NAMES
```

If you do not provide a name for the container during its creation, then Podman generates a random string name for the container. It is important to define a unique name to facilitate the identification of your containers when managing their lifecycle.

You can assign a name to your containers by adding the `--name` flag to the `podman run` command.

```
[user@host ~]$ podman run --name podman_rhel9 \
registry.redhat.io/rhel9/rhel-guest-image:9.4 echo 'Red Hat'
Red Hat
```

Podman can identify the containers either by the *Universal Unique Identifier (UUID)* short identifier, which is composed of twelve alphanumeric characters, or by the UUID long identifier, which is composed of 64 alphanumeric characters, as shown in the example.

```
[user@host ~]$ podman ps --all
CONTAINER ID  IMAGE                                COMMAND                                CREATED
STATUS        PORTS        NAMES
20236410bcef  registry.redhat.io/rhel9/rhel-guest-image:9.4  echo Red Hat  1
second ago    Exited (0) 1 second ago                podman_rhel9
```

In this example, `20236410bcef` is the container's ID, or UUID short identifier. Additionally, `podman_rhel9` is listed as the container's name.

If you want to retrieve the UUID long container ID, then you can add the `--format=json` flag to the `podman ps --all` command.

```
[user@host ~]$ podman ps --all --format=json
{
  "AutoRemove": false,
  "Command": [
    "echo",
    "Red Hat"
  ],
  ...output omitted...
  "Id": "2024...b0b2", ❶
  "Image": "registry.redhat.io/rhel9/rhel-guest-image:9.4",
  ...output omitted...
```

❶ UUID long container ID.

You can retrieve information about a container in either the JSON format or as a Go template.

Exposing Containers

Many applications, such as web servers or databases, run indefinitely to wait for connections. Therefore, the containers for these applications must run indefinitely. At the same time, it is usually necessary for these applications to be accessed externally through a network protocol.

You can use the `-p` option to map a port in your local machine to a port inside the container. This way, the traffic in your local port is forwarded to the port inside the container, which enables you to access the application from your computer.

The following example creates a new container that runs an Apache HTTP server by mapping the 8080 port in your local machine to the 8080 port inside the container.

```
[user@host ~]$ podman run -p 8080:8080 \
  registry.access.redhat.com/ubi9/httpd-24:latest
...output omitted...
[Thu Jun 18 12:58:57.048491 2024] [ssl:warn] [pid 1:tid 140259.33613248] AH01909:
10.0.2.100:8443:0 server certificate does NOT include an ID which matches the
server name
[Thu Jun 18 12:58:57.048899 2024] [:notice] [pid 1:tid 140259.33613248]
ModSecurity for Apache/2.9.2 (http://www.modsecurity.org/) configured.
...output omitted...
[Thu Jun 18 12:58:57.136272 2024] [mpm_event:notice] [pid 1:tid 140259.33613248]
AH00489: Apache/2.4.37 (Red Hat Enterprise Linux) OpenSSL/1.1.1k configured --
resuming normal operations
[Thu Jun 18 12:58:57.136332 2024] [core:notice] [pid 1:tid 140259.33613248]
AH00094: Command line: 'httpd -D FOREGROUND'
```

You can access the HTTP server at `localhost:8080`.

If you want the container to run in the background, to avoid the terminal being blocked, then you can use the `-d` option.

```
[user@host ~]$ podman run -d -p 8080:8080 \
  registry.access.redhat.com/ubi9/httpd-24:latest
b7eb467781106e4f416ba79cede91152239bfc74f6a570c6d70baa4c64fa636a
```

Using Environment Variables

Environment variables are variables used in your applications that are set outside of the program. The operating system or the environment where the application runs provides the value of the variable. You can access the environment variable in your application at runtime.

Environment variables are a useful and safe way of injecting environment-specific configuration values into your application. For example, your application might use a database hostname that is different for each application environment, such as the `database.local`, `database.stage`, or `database.test` hostnames.


You can pass environment variables to a container by using the `-e` option. In the following example, an environment variable called `NAME` with the value `Red Hat` is passed. Then, the environment variable is printed by using the `printenv` command inside the container.

```
[user@host ~]$ podman run -e NAME='Red Hat' \
registry.redhat.io/rhel9/rhel-guest-image:9.4 printenv NAME
Red Hat
```


Podman Desktop

Podman Desktop is a graphical user interface that you can use to manage and interact with containers in local environments. It uses the Podman engine by default, and supports other container engines as well, such as Docker.


On launch, Podman Desktop displays a dashboard with information about the status of the Podman engine. The dashboard might display warnings or errors, for example, if the Podman engine is not installed, or if the Docker compatibility is not fully set up.



With Podman Desktop, you can perform many of the tasks that you can do with the `podman` CLI, such as pulling images and creating containers. For example, you can create, list, and run containers from the **Containers** section, as the following image shows:



Similarly, you can pull, list images, and create containers from those images in the **Images** section.



Podman Desktop is an addition to the `podman` CLI, rather than a replacement. For more advanced commands or options, the CLI is required. Still, Podman Desktop can be useful for users who prefer graphical environments for common tasks, and for beginners who are learning about containers.

Podman Desktop is modular and extensible. You can use and create extensions that provide additional capabilities. For example, with the *Red Hat OpenShift* extension, Podman Desktop can deploy containers to Red Hat OpenShift Container Platform.

Podman Desktop is available for Linux, MacOS, and Windows. For specific installation instructions, refer to the Podman Desktop documentation.



Note

Podman Desktop is not shipped with Red Hat Enterprise Linux. However, Red Hat provides a free download at <https://developers.redhat.com/products/podman-desktop/overview>.



References

Podman Documentation

<https://docs.podman.io/en/latest/>

Formatting Output with Podman

<https://docs.podman.io/en/latest/markdown/podman-ps.1.html>

Podman Desktop Documentation

<https://podman-desktop.io/docs/intro>

► Guided Exercise

Creating Containers with Podman

Use Podman to pull container images from a registry and to use them to run containers.

Outcomes

You should be able to run containers in Podman by using the `podman run` command.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

```
[student@workstation ~]$ lab start containers-creating
```

Instructions

- 1. Use the `registry.lab.example.com/ubi9/ubi` image to create a new container that prints the `Hello Red Hat` text.

- 1.1. Log in to the classroom registry. Use `admin` as the username, and `redhat321` as the password.

```
[student@workstation ~]$ podman login registry.lab.example.com --tls-verify=false
Username: admin
Password: redhat321
Login Succeeded!
```

- 1.2. Use the `podman pull` command to fetch the image from the registry.

```
[student@workstation ~]$ podman pull registry.lab.example.com/ubi9/ubi \
--tls-verify=false
Trying to pull registry.lab.example.com/ubi9/ubi:latest...
...output omitted...
Getting image source signatures
Copying blob 1153e061da4e done
Copying config 8d2a8803cf done
Writing manifest to image destination
8d2a8803cfca17a81eb9412e1f33ae1c6fe3797553e9b819899dc03f1657cf12
```

- 1.3. Use the `podman images` command to verify that the image is available locally.

```
[student@workstation ~]$ podman images
REPOSITORY                                TAG          IMAGE ID      CREATED      SIZE
registry.lab.example.com/ubi9/ubi         latest       8d2a8803cfca  3 months ago  219 MB
```

- 1.4. In a command line terminal, use `podman run` to create a new container. Provide an `echo` command to be executed inside the container.

```
[student@workstation ~]$ podman run --rm registry.lab.example.com/ubi9/ubi \
echo 'Hello Red Hat'
Hello Red Hat
```

- 1.5. Verify that the container is not running after the execution finishes.

```
[student@workstation ~]$ podman ps
```

CONTAINER ID	IMAGE	COMMAND	CREATED	STATUS	PORTS	NAMES

- ▶ 2. Explore setting and printing environment variables by using the `registry.lab.example.com/ubi9/ubi` container image.

- 2.1. Use the `-e` option in the `podman run` command to set environment variables. Use the `printenv` command that is packaged in the container image to print the values of the environment variables.

```
[student@workstation ~]$ podman run --rm -e GREET=Hello -e NAME='Red Hat' \
registry.lab.example.com/ubi9/ubi printenv GREET NAME
Hello
Red Hat
```

- 2.2. Verify that the container is not running after the execution finishes.

```
[student@workstation ~]$ podman ps
```

CONTAINER ID	IMAGE	COMMAND	CREATED	STATUS	PORTS	NAMES

- ▶ 3. Use the `registry.lab.example.com/rhel9/httpd-24` image to run a new container that creates an Apache HTTP server.

- 3.1. Use the `podman pull` command to fetch the image from the registry.

```
[student@workstation ~]$ podman pull registry.lab.example.com/rhel9/httpd-24 \
--tls-verify=false
Trying to pull registry.lab.example.com/rhel9/httpd-24:latest...
...output omitted...
Copying blob 358ca3c2eaff done
Copying blob e584cd196457 done
Copying config a42923e055 done
Writing manifest to image destination
a42923e05589d2d10c25e87736bc3b9537c38c2a6a87f37e9a71e288b01738a6
```

- 3.2. Use the `podman run` command to start the `httpd` container. Use the `-p` option in the `podman run` command to redirect traffic from port `8080` on your machine to port `8080` inside the container.

```
[student@workstation ~]$ podman run --rm -p 8080:8080 \
  registry.lab.example.com/rhel9/httpd-24
=> sourcing 10-set-mpm.sh ...
=> sourcing 20-copy-config.sh ...
=> sourcing 40-ssl-certs.sh ...
--> Generating SSL key pair for httpd...
...output omitted...
[Thu Jun 06 00:02:06.591391 2024] [mpm_event:notice] [pid 1:tid 1] AH00489:
  Apache/2.4.57 (Red Hat Enterprise Linux) OpenSSL/3.0.7 configured -- resuming
  normal operations
[Thu Jun 06 00:02:06.591430 2024] [core:notice] [pid 1:tid 1] AH00094: Command
  line: 'httpd -D FOREGROUND'
```

The container keeps running and the logs are displayed in your terminal.

- 3.3. In a web browser, navigate to `http://localhost:8080`. Verify that the HTTP server is running at the 8080 port.
- 3.4. Go back to your command-line terminal. Press `Ctrl+C` to stop the container.

```
...output omitted...
^C
[Thu Jun 06 00:04:13.467695 2024] [mpm_event:notice] [pid 1:tid 1] AH00491: caught
  SIGTERM, shutting down
```

- 3.5. Create the container again by adding the `-d` option. The container runs in the background.


```
[student@workstation ~]$ podman run --rm -d \
  -p 8080:8080 registry.lab.example.com/rhel9/httpd-24
08ac7ce948feb65f76f82305d556843a898b60200f5b4f37d1ecb8ea381199ff
```

- 3.6. Verify that the container is running in the background.


```
[student@workstation ~]$ podman ps
```

CONTAINER ID	IMAGE	COMMAND
CREATED	STATUS	PORTS
		NAMES
08ac7ce948fe	registry.lab.example.com/rhel9/httpd-24:latest	/usr/bin/run-http...
21 seconds ago	Up 22 seconds	0.0.0.0:8080->8080/tcp charming_tharp


- ▶ 4. Use Podman Desktop to list images and containers.
 - 4.1. Open Podman Desktop. Click **Activities** at the upper-left corner of the top bar, and then click the **Podman Desktop** icon.




- 4.2. Click **Containers** in the Podman Desktop navigation panel. The list includes the containers that you created in the exercise. Verify that the Apache HTTP container is running.



- 4.3. Click **Images**. The list of local images displays the ubi9/ubi and the rhel9/httpd-24 images.



5. Use Podman Desktop to start another Apache HTTP container that maps its port to port 8090 in the workstation.
 - 5.1. Click the arrow (▶) of the rhel9/httpd-24 image to run a new container based on this image.
 - 5.2. In the container creation form, enter the following values:
 - Container name: **desktop-test**
 - Local port for 8080/tcp: **8090**
 - Local port for 8443/tcp: **8443**




5.3. Click **Start Container**.



Note

You can safely ignore the SELinux security warning that displays at the top of the desktop.

- 5.4. Navigate to the containers menu and verify that the **desktop-test** container is running.
- 5.5. In the **desktop-test** container, click the vertical ellipsis icon (⋮) on the right, select **Open Browser**, and then click **Yes** in the confirmation dialog box.



- 5.6. Verify that the browser can access the HTTP server at `http://localhost:8090`.
- 5.7. Return to Podman Desktop. In the `desktop-test` container, click the **Delete Container** icon to delete the container, and then close Podman Desktop.

Finish

On the workstation machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish containers-creating
```

Container Image Registries

Objectives

- Authenticate to and manage container image registries.

Container Registries

A container image is a packaged version of your application, with all the dependencies that are necessary for the application to run. You can use image registries to store container images to later share them in a controlled manner. Some examples of image registries include Quay.io, Red Hat Registry, Docker Hub, and Amazon ECR.

For example, consider the following Podman command.

```
[user@host ~]$ podman pull registry.redhat.io/ubi9/ubi:9.1
Trying to pull registry.redhat.io/ubi9/ubi:9.1...
Getting image source signatures
...output omitted...
Writing manifest to image destination
Storing signatures
3434...8f6b
```

The `registry.redhat.io/ubi9` image is stored in the Red Hat Registry, because the container name uses the `registry.redhat.io` host.

Red Hat Registry

Red Hat distributes container images by using two registries:

- `registry.access.redhat.com`: requires no authentication
- `registry.redhat.io`: requires authentication

However, Red Hat provides a centralized searching utility for both registries: the Red Hat Ecosystem Catalog, available at <https://catalog.redhat.com/>. You can use the Red Hat Ecosystem Catalog to search for images and get technical details about them. Navigate to <https://catalog.redhat.com/software/containers/explore> to search for container images.




Figure 13.10: The Red Hat Ecosystem Catalog

The container image details page provides relevant information about the container image, such as the Containerfile used to create the image, the packages installed within the image, or a security scanning. You can also change the image version by selecting a specific *tag*.




Figure 13.11: The Node.js 16 image based on RHEL 9

Useful Container Images

The following table introduces some useful container images.

Image	Provides	Description
registry.access.redhat.com/ubi9	Universal Base Image (UBI), Version 9	A base image to create other images that is based on RHEL 9.

Image	Provides	Description
<code>registry.access.redhat.com/ubi9/python-312</code>	Python 3.12	A UBI-based image with the Python 3.12 runtime.
<code>registry.access.redhat.com/ubi9/nodejs-18</code>	Node.js 18	A UBI-based image with the Node.js 18 runtime.
<code>registry.access.redhat.com/ubi9/go-toolset</code>	Go Toolset	A UBI-based image with the Go runtime. The Go version depends on the image tag.

Red Hat UBIs are *Open Container Initiative (OCI)* compliant enterprise grade container images that provide the base operating system layer for your containerized applications. UBIs include a subset of Red Hat Enterprise Linux (RHEL) components. UBIs can additionally provide a set of pre-built language runtimes. UBIs are freely distributable, and you can use UBIs on both Red Hat and non-Red Hat platforms or container registries.

You do not need a Red Hat subscription to use or distribute UBI-based images. However, Red Hat only provides full support for containers that are built on UBI if the containers are deployed to a Red Hat platform, such as Red Hat OpenShift Container Platform (RHOCP) or RHEL.

Quay.io

Although the Red Hat Registry only stores images from Red Hat and certified providers, you can use the Quay.io registry to store your custom images. Storing public images in Quay.io is free, and paying customers receive further benefits, such as private repositories. Developers can also deploy an on-premise Quay instance, which you can use to set up an image registry on your infrastructure.

To log in to Quay.io, you can use your Red Hat developer account.




Figure 13.12: The Quay.io welcome page

Manage Registries with Podman

When you pull a container image, you provide a number of details. For example, the `registry.access.redhat.com/ubi9/nodejs-18:latest` image name consists of the following information:

- Registry URL: `registry.access.redhat.com`
- User or organization: `ubi9`
- Image repository: `nodejs-18`
- Image tag: `latest`

Developers can specify a shorter, unqualified name, which omits the registry URL. For example, you might shorten the `registry.redhat.io/ubi9/nodejs-18:latest` to `ubi9/nodejs-18:latest`.

```
[user@host ~]$ podman pull ubi9/nodejs-18
```

If you do not provide the registry URL, then Podman uses the `/etc/containers/registries.conf` file to search other container registries that might contain the image name. This file contains registries that Podman searches to find the image, in order of preference.

For example, with the following configuration, Podman searches the Red Hat Registry first. If the image is not found in the Red Hat Registry, then Podman searches in the Docker Hub registry.

```
unqualified-search-registries == ['registry.redhat.io', 'docker.io']
```

You can also block a registry. For example, the following configuration blocks pulling from Docker Hub.

```
[[registry]]
location="docker.io"
blocked=true
```

Some systems do not have a `/etc/containers/registries.conf` file, such as Microsoft Windows. In such cases, the `registries.conf` file might exist in a different location.

For example, on Microsoft Windows, execute the `podman machine ssh` command to connect to the Linux-based virtual machine that starts your containers. In the virtual machine, you can find the `/etc/containers/registries.conf` file:

```
[user@host ~]$ podman machine ssh
[user@DESKTOP-AA1A111 ~]$ ls /etc/containers/containers.conf
/etc/containers/containers.conf
```

See the references section for more details about the `registries.conf` file.

Red Hat recommends that you always use a fully qualified container image name to avoid duplicate container images in multiple container registries. For example, depending on your Podman configuration, the `rhel9/rhel-guest-image:9.4` container image might resolve to a potentially unsupported or malicious `docker.io/library/rhel9/rhel:9.4` container image.

Manage Registries with Skopeo

Skopeo is a command-line tool for working with container images. Developers can use *Skopeo* for various tasks:

- Inspect remote container images.
- Copy a container image between registries.
- Sign an image with OpenPGP keys.
- Convert image formats, for example from **Docker** to the **OCI** format.

Skopeo can inspect remote images or transfer images between registries without using local storage. The *skopeo* command uses the `transport:image` format, such as `docker://remote_image`, `dir:path`, or `oci:path:tag`.

Use the `skopeo inspect` command to read image metadata.

```
[user@host ~]$ skopeo inspect \
  docker://registry.access.redhat.com/ubi9/nodejs-18
{
  "Name": "registry.access.redhat.com/ubi9/nodejs-18",
  "Digest": "sha256:741b...22e0",
  "RepoTags": [
    ...output omitted...
```

Use the `skopeo copy` command to copy images between registries. The following example copies the `registry.access.redhat.com/ubi9/nodejs-18:latest` image into the `quay.io/myuser/nodejs-18` Quay.io repository.

```
[user@host ~]$ skopeo copy \
  docker://registry.access.redhat.com/ubi9/nodejs-18 \
  docker://quay.io/myuser/nodejs-18
Getting image source signatures
...output omitted...
```

The following example changes the transport format to download an image into a local directory.

```
[user@host ~]$ skopeo copy \
  docker://registry.access.redhat.com/ubi9/nodejs-18 \
  dir:/var/lib/images/nodejs-18
Getting image source signatures
...output omitted...
```

See the references section for more details about *Skopeo*.

Manage Registry Credentials with Podman

Some registries require users to authenticate, such as the `registry.redhat.io` registry.

```
[user@host ~]$ podman pull registry.redhat.io/rhel9/nginx-120
Trying to pull registry.redhat.io/rhel9/nginx-120:latest...
Error: initializing source docker://registry.redhat.io/ubi9/
httpd-24:latest: unable to retrieve auth token: invalid username/password:
unauthorized: Please login to the Red Hat Registry using your Customer Portal
credentials. Further instructions can be found here: https://access.redhat.com/
RegistryAuthentication
```

You might choose a different image that does not require authentication, such as the UBI 9 image from the `registry.access.redhat.com` registry:

```
[user@host ~]$ podman pull registry.access.redhat.com/ubi9:latest
Trying to pull registry.access.redhat.com/ubi9:latest...
Getting image source signatures
Checking if image destination supports signatures
...output omitted...
```

Alternatively, authenticate your calls by executing the `podman login` command.

```
[user@host ~]$ podman login registry.redhat.io
Username: YOUR_USER
Password: YOUR_PASSWORD
Login Succeeded!
[user@host ~]$ podman pull registry.redhat.io/rhel9/nginx-120
Trying to pull registry.redhat.io/rhel9/nginx-120:latest...
Getting image source signatures
...output omitted...
```

Podman stores the credentials in the `${XDG_RUNTIME_DIR}/containers/auth.json` file, where the `${XDG_RUNTIME_DIR}` refers to a directory specific to the current user. The credentials are encoded in the base64 format:

```
[user@host ~]$ cat ${XDG_RUNTIME_DIR}/containers/auth.json
{
  "auths": {
    "registry.redhat.io": {
      "auth": "dXNlcjpodW50ZXIy"
    }
  }
}
[user@host ~]$ echo -n dXNlcjpodW50ZXIy | base64 -d
user:hunter2
```



Note

For security reasons, the `podman login` command does not show your password in the interactive session. Although you do not see what you are typing, Podman registers every key stroke. Press Enter when you have typed your full password in the interactive session to initiate the login.

Skopeo uses the same `${XDG_RUNTIME_DIR}/containers/auth.json` file to access authentication details for each registry.



References

`skopeo(1)` and `skopeo-copy(1)` man pages

Red Hat - How to Manage Linux Container Registries

<https://www.redhat.com/sysadmin/manage-container-registries>

► Guided Exercise

Container Image Registries

Use Podman to pull container images from a registry, and use the images to run containers.

Outcomes

- Configure a container image registry and create a container from an existing image.
- Delete containers and images.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start containers-basics
```

Instructions

- 1. Log in to the servera machine as the student user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
```

- 2. Install the `container-tools` metapackage. Use `student` as the password.

```
[student@servera ~]$ sudo dnf install container-tools
[sudo] password for student: student
...output omitted...
Is this ok [y/N]: y
...output omitted...
Complete!
```

- 3. Configure the `registry.lab.example.com` classroom registry in your home directory. Log in to the container registry with `admin` as the user and `redhat321` as the password.

- 3.1. Create the `/home/student/.config/containers` directory.

```
[student@servera ~]$ mkdir -p /home/student/.config/containers
```

- 3.2. Create the `/home/student/.config/containers/registries.conf` file with the following contents:

```
unqualified-search-registries = ['registry.lab.example.com']

[[registry]]
location = "registry.lab.example.com"
insecure = true
blocked = false
```

3.3. Verify that the classroom registry is added.

```
[student@servera ~]$ podman info
...output omitted...
registries:
  registry.lab.example.com:
    Blocked: false
    Insecure: true
    Location: registry.lab.example.com
    MirrorByDigestOnly: false
    Mirrors: null
    Prefix: registry.lab.example.com
  search:
    - registry.lab.example.com
...output omitted...
```

3.4. Log in to the classroom registry.

```
[student@servera ~]$ podman login registry.lab.example.com
Username: admin
Password: redhat321
Login Succeeded!
```

- 4. Run the `python312` container in detached mode from an image with the `python 3.12` package and based on the `ubi9` image. The image is hosted on a remote registry.

4.1. Search for a `python-312` container in the `registry.lab.example.com` registry.

```
[student@servera ~]$ podman search registry.lab.example.com/
NAME                                DESCRIPTION
...output omitted...
registry.lab.example.com/rhel9/mariadb-105
registry.lab.example.com/rhel9/httpd-24
registry.lab.example.com/ubi9/python-312
...output omitted...
```

4.2. Use the `skopeo` command to inspect the image.

```
[student@servera ~]$ skopeo inspect \
docker://registry.lab.example.com/ubi9/python-312
...output omitted...
"description": "Python 3.12 available as container is a base platform for
building and running various Python 3.12 applications and frameworks.
...output omitted..."
```

4.3. Pull the python-312 container image.

```
[student@servera ~]$ podman pull registry.lab.example.com/ubi9/python-312
Trying to pull registry.lab.example.com/ubi9/python-312:latest...
...output omitted...
7d85f178ae10929ae2fbe742a70d91b8ef432dbe84ff55d0af9f05d9506e9890
```

4.4. Verify that the container is downloaded to the local image repository.

```
[student@servera ~]$ podman images
```

REPOSITORY	TAG	IMAGE ID	CREATED	SIZE
registry.lab.example.com/ubi9/python-312	latest	7d85f178ae10	2 weeks ago	1.08 GB

4.5. Start the python312 container.

```
[student@servera ~]$ podman run -d --name python312 \
registry.lab.example.com/ubi9/python-312 sleep infinity
afccc88bd61321b7d7466d71fd692b7bafb63c3e160435fd266e1d50d6263989
```

4.6. Verify that the container was created.

```
[student@servera ~]$ podman ps
```

CONTAINER ID	IMAGE	COMMAND
afccc88bd613	registry.lab.example.com/ubi9/python-312:latest	sleep infinity
seconds ago	Up 21 seconds	python312

► 5. Delete the container and image. Return to workstation.

5.1. Stop the container.

```
[student@servera ~]$ podman stop python312
...output omitted...
python312
```

5.2. Remove the container.

```
[student@servera ~]$ podman rm python312
python312
```

5.3. Remove the container image.

```
[student@servera ~]$ podman rmi \
registry.lab.example.com/ubi9/python-312:latest
Untagged: registry.lab.example.com/ubi9/python-312:latest
Deleted: 7d85...9890
```

5.4. Return to workstation.


```
[student@servera ~]$ exit  
logout  
Connection to servera closed.
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish containers-basics
```

Managing the Container Lifecycle

Objectives

- Manage the lifecycle of a container from creation to deletion.

Container Lifecycle

Podman provides a set of subcommands to create and manage containers. You can use those subcommands to manage the container and container image lifecycle.

The following figure shows the most commonly used Podman subcommands that change the container and image state:




Figure 13.13: Podman lifecycle commands

Podman also provides a set of subcommands to obtain information about running and stopped containers.

You can use these subcommands to extract information from containers and images for debugging, updating, or reporting purposes. The following figure shows the most commonly used subcommands that query information from containers and images:




Figure 13.14: Podman query commands

This lecture covers the basic operations that you can use to manage containers. Commands that are explained in this lecture accept either a container ID or the container name.

Listing Containers

You can list running containers with the `podman ps` command.

```
[user@host ~]$ podman ps
CONTAINER ID  IMAGE           COMMAND          CREATED  STATUS  PORTS          NAMES
0ae7be593698  ...server      /bin/sh -c python...  ...ago  Up...   ...8000/tcp    httpd
c42e7dca12d9  ...helloworld  /bin/sh -c nginx...  ...ago  Up...   ...8080/tcp    nginx
```

Each row describes information about the container, such as the image that is used to start the container, the command executed when the container started, and the container uptime.

You can include stopped containers in the output by adding the `--all` or `-a` flag to the `podman ps` command.

```
[user@host ~]$ podman ps --all
```

CONTAINER ID	IMAGE	COMMAND	CREATED	STATUS	PORTS
0ae7be593698	...server	/bin/sh -c python...	...ago	Up...	...8000/tcp
bd5ada1b6321	...httpd-24	/usr/bin/run-httpd...	...ago	Exited...	...8080/tcp
c42e7dca12d9	...helloworld	/bin/sh -c nginx	...ago	Up...	...8080/tcp

Inspecting Containers

In the context of Podman, inspecting a container means retrieving the full information of the container. The `podman inspect` command returns a JSON array with information about different aspects of the container, such as networking settings, CPU usage, environment variables, status, port mapping, or volumes. The following snippet is a sample output of the `podman inspect` command.

```
[user@host ~]$ podman inspect 7763097d11ab
```

```
[
  {
    "Id": "7763...cbc0",
    "Created": "2022-05-04T10:00:32.988377257-03:00",
    "Path": "container-entrypoint",
    "Args": [
      "/usr/bin/run-httpd"
    ],
    "State": {
      "OciVersion": "1.0.2-dev",
      "Status": "running",
      "Running": true,
      "Paused": false,
      "Restarting": false,
      "OOMKilled": false,
      "Dead": false,
      "Pid": 9746,
      ...output omitted...
      "Image": "d2b9...fa0a",
      "ImageName": "registry.access.redhat.com/ubi8/httpd-24:latest",
      "Rootfs": "",
      ...output omitted...
      "Env": [
        "PATH=/opt/app-root/src/bin:/opt/app-root/bin:/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin",
        "TERM=xterm",
        "container=oci",
        "HTTPD_VERSION=2.4",
        ...output omitted...
        "CpuCount": 0,
        "CpuPercent": 0,
        "IOMaximumIOps": 0,
        "IOMaximumBandwidth": 0,
        "CgroupConf": null
      ]
    }
  }
]
```

```
    }
  }
]
```

The previous JSON output omits most fields. Because the full JSON output is long, you might find it difficult to find the specific fields that you want to retrieve. You can pass a `--format` flag to the `podman inspect` command to filter the command output. The `--format` flag expects a Go template expression as a parameter, which you can use to select specific fields in the JSON output.

Go template expressions use annotations, which are delimited by double curly braces (`{{` and `}}`), to refer to elements, such as fields or keys, in a data structure. The data structure elements are separated by a dot (`.`) and must start in uppercase.

In the following command, the `Status` field of the `State` object is retrieved for a container called `redhat`.

```
[user@host ~]$ podman inspect \
  --format='{{.State.Status}}' redhat
running
```

In the preceding example, the `podman inspect` command uses the supplied Go template to navigate the JSON array of the `redhat` container. The template returns the `Status` field value of the `State` object from the JSON array.

Refer to the references section for more information about the Go templating language.

Stopping Containers Gracefully

You can stop a container gracefully by using the `podman stop` command. When you execute the `podman stop` command, Podman sends a `SIGTERM` signal to the container. Processes use the `SIGTERM` signal to implement clean-up procedures before stopping.

The following command stops a container with a container ID `1b982aeb75dd`.

```
[user@host ~]$ podman stop 1b982aeb75dd
1b982aeb75dd
```

You can stop all the running containers by using the `--all` or `-a` flag. In the following example, the command stops three containers.

```
[user@host ~]$ podman stop --all
4aea...0c2a
6b18...54ea
7763...cbc0
```

If a container does not respond to the `SIGTERM` signal, then Podman sends a `SIGKILL` signal to forcefully stop the container. Podman waits 10 seconds by default before sending the `SIGKILL` signal. You can change the default behavior by using the `--time` flag.

```
[user@host ~]$ podman stop --time=100
```

In the previous example, Podman gives the container a grace period of 100 seconds before sending the killing signal.

Stopping a container differs from removing a container. Stopped containers are present on the host machine and can be listed by using the `podman ps --all` command.

Additionally, when the containerized process finishes, the container enters the exited state. Such a process might finish for a number of reasons, such as an error, OOM (out-of-memory) state, or successfully finishing. Podman lists exited containers with other stopped containers.

Stopping Containers Forcefully

You can send the `SIGKILL` signal to the container by using the `podman kill` command. In the following example, a container called `httpd` is stopped forcefully.

```
[user@host ~]$ podman kill httpd
httpd
```

Pausing Containers

Both `podman stop` and `podman kill` commands eventually send a `SIGKILL` signal to the container. The `podman pause` command suspends all processes in the container by sending the `SIGSTOP` signal.

```
[user@host ~]$ podman pause 4f2038c05b8c
4f2038c05b8c
```

The `podman unpause` command resumes a paused container.

```
[user@host ~]$ podman unpause 4f2038c05b8c
4f2038c05b8c
```

Restarting Containers

Execute the `podman restart` command to restart a running container. You can also use the command to start stopped containers.

The following command restarts a container called `nginx`.

```
[user@host ~]$ podman restart nginx
1b98...75dd
```

Removing Containers

Use the `podman rm` command to remove a stopped container. The following command removes a stopped container with the container ID `c58cfd4b90df`.

```
[user@host ~]$ podman rm c58cfd4b90df
c58c...90df
```

You cannot remove running containers by default. You must stop the running container first and then remove it. The following command tries to remove a running container.

```
[user@host ~]$ podman rm c58cfd4b90df
Error: cannot remove container c58c...90df as it is running - running or paused
containers cannot be removed without force: container state improper
```

You can add the `--force` (or `-f`) flag to remove the container forcefully.

```
[user@host ~]$ podman rm c58cfd4b90df --force
c58c...90df
```

You can also add the `--all` (or `-a`) flag to remove all stopped containers. This flag fails to remove running containers. The following command removes two containers.

```
[user@host ~]$ podman rm --all
6b18...54ea
6c0d...a6fb
```

You can combine the `--force` and `--all` flags to remove all containers, including running containers.

Container Persistent Storage

By default, when you run a container, all of the content uses the container-based image. Given the ephemeral nature of container images, all of the new data that the user or the application writes is lost after removing the container.

To make data persistent, you can use host file-system content in the container with the `--volume` (`-v`) option. You must consider file-system level permissions when you use this volume type in a container.

In the MariaDB container image, the `mysql` user must own the `/var/lib/mysql` directory, the same as if MariaDB was running on the host machine. The directory to mount into the container must have `mysql` as the user and group owner (or the UID and GID of the `mysql` user, if MariaDB is not installed on the host machine). If you run a container as the `root` user, then the UIDs and GIDs on your host machine match the UIDs and GIDs inside the container.

The UID and GID matching configuration does not occur the same way in a rootless container. In a rootless container, the user has root access from within the container, because Podman launches a container inside the user namespace.

You can use the `podman unshare` command to run a command inside the user namespace. To obtain the UID mapping for your user namespace, use the `podman unshare cat` command.

```
[user@host ~]$ podman unshare cat /proc/self/uid_map
0      1000      1
1      100000    65536
[user@host ~]$ podman unshare cat /proc/self/gid_map
0      1000      1
1      100000    65536
```

The preceding output shows that in the container, the root user (UID and GID of 0) maps to your user (UID and GID of 1000) on the host machine. In the container, the UID and GID of 1 maps to the UID and GID of 100000 on the host machine. Every UID and GID after 1 increments by 1. For

example, the UID and GID of 30 inside a container maps to the UID and GID of 100029 on the host machine.

You use the `podman exec` command to view the `mysql` user UID and GID inside the container that is running with ephemeral storage.

```
[user@host ~]$ podman exec -it db01 grep mysql /etc/passwd
mysql:x:27:27:MySQL Server:/var/lib/mysql:/sbin/nologin
```

You decide to mount the `/home/user/db_data` directory into the `db01` container to provide persistent storage on the `/var/lib/mysql` directory of the container. You then create the `/home/user/db_data` directory, and use the `podman unshare` command to set the user namespace UID and GID of 27 as the owner of the directory.

```
[user@host ~]$ mkdir /home/user/db_data
[user@host ~]$ podman unshare chown 27:27 /home/user/db_data
```

The UID and GID of 27 in the container maps to the UID and GID of 100026 on the host machine. You can verify the mapping by viewing the ownership of the `/home/user/db_data` directory with the `ls` command.

```
[user@host ~]$ ls -l /home/user/
total 0
drwxrwxr-x. 3 100026 100026 18 May  5 14:37 db_data
...output omitted...
```

Now that the correct file-system level permissions are set, you use the `podman run` command `-v` option to mount the directory.

```
[user@host ~]$ podman run -d --name db01 \
-e MYSQL_USER=student \
-e MYSQL_PASSWORD=student \
-e MYSQL_DATABASE=dev_data \
-e MYSQL_ROOT_PASSWORD=redhat \
-v /home/user/db_data:/var/lib/mysql \
registry.lab.example.com/rhel8/mariadb-105
```

You notice that the `db01` container is not running.

```
[user@host ~]$ podman ps -a
```

CONTAINER ID	IMAGE	COMMAND
dfdc20cf9a7e	registry.lab.example.com/rhel8/mariadb-105:latest	run-mysqld
29 seconds ago	Exited (1) 29 seconds ago	db01

The `podman container logs` command shows a permission error for the `/var/lib/mysql/db_data` directory.


```
[user@host ~]$ podman container logs db01
...output omitted...
---> 16:41:25      Initializing database ...
---> 16:41:25      Running mysql_install_db ...
mkdir: cannot create directory '/var/lib/mysql/db_data': Permission denied
Fatal error Can't create database directory '/var/lib/mysql/db_data'
```

This error happens because of the incorrect SELinux context that is set on the `/home/user/db_data` directory on the host machine.

SELinux Contexts for Container Storage

You must set the `container_file_t` SELinux context type before you can mount the directory as persistent storage to a container. If the directory does not have the `container_file_t` SELinux context, then the container cannot access the directory. You can append the `Z` option to the argument of the `podman run` command `-v` option to automatically set the SELinux context on the directory.

So you use the `podman run -v /home/user/db_data:/var/lib/mysql:Z` command to set the SELinux context for the `/home/user/db_data` directory when you mount it as persistent storage for the `/var/lib/mysql` directory.

```
[user@host ~]$ podman run -d --name db01 \
-e MYSQL_USER=student \
-e MYSQL_PASSWORD=student \
-e MYSQL_DATABASE=dev_data \
-e MYSQL_ROOT_PASSWORD=redhat \
-v /home/user/db_data:/var/lib/mysql:Z \
registry.lab.example.com/rhel8/mariadb-105
```

You then verify that the correct SELinux context is set on the `/home/user/db_data` directory with the `ls` command `-Z` option.

```
[user@host ~]$ ls -Z /home/user/
system_u:object_r:container_file_t:s0:c81,c1009 db_data
...output omitted...
```

Start a Containerized Service on Boot

In a traditional environment, administrators configure applications such as web servers or databases to start at boot time, and run indefinitely as a `systemd` service. `Systemd` is a system and service management tool for Linux operating systems. `Systemd` uses service unit files to start and stop applications, or to enable them to start at boot time. Typically, an administrator manages these applications with the `systemctl` command. As a regular user, you can create `systemd` unit files to manage your rootless containers. You can use this configuration to manage your container as a regular `systemd` service with the `systemctl` command.

Generate a Systemd Unit File

To create a `systemd` unit file for a specified service container, use the `podman generate systemd` command. Use the `--name` option to specify the container's name. The `--files` option generates files instead of printing to standard output (STDOUT).

```
[user@host ~]$ podman generate systemd --name web --files
/home/user/container-web.service
```

The preceding command creates the `container-web.service` unit file for a container called `web`. After reviewing and adapting the generated service configuration to your requirements, store the unit file in the user's `systemd` configuration directory (`~/.config/systemd/user/`).

After adding or modifying unit files, you must use the `systemctl` command to reload the `systemd` configuration.

```
[user@host ~]$ systemctl --user daemon-reload
```

Managing the Containerized Service

To manage a containerized service, use the `systemctl` command.

```
[user@host ~]$ systemctl --user [start, stop, status, enable, disable]
container-web.service
```

When you use the `--user` option, by default, `systemd` starts the service when you log in, and stops it when you log out. You can start your enabled services at the operating system boot, and stop them on shutdown, by running the `loginctl enable-linger` command.

```
[user@host ~]$ loginctl enable-linger
```

To revert the operation, use the `loginctl disable-linger` command.



References

`podman-inspect(1)`, `podman-stop(1)`, `podman-restart(1)`, `podman-rm(1)`, `podman-run(1)`, and `podman-unshare(1)` man pages

Go Templates Package

<https://pkg.go.dev/text/template>

For more information, refer to the *Porting Containers to systemd Using Podman* chapter in the *Red Hat Enterprise Linux 9 Building, Running, and Managing Containers* guide at

https://docs.redhat.com/en/documentation/red_hat_enterprise_linux/9/html-single/building_running_and_managing_containers/index#assembly_porting-containers-to-systemd-using-podman_building-running-and-managing-containers

► Guided Exercise

Managing the Container Lifecycle

Manage the lifecycle of a container that runs an Apache HTTP Server.

Outcomes

You should be able to:

- Get detailed information about a container.
- Stop containers.
- Restart a stopped container.
- Delete containers.
- Create and mount persistent storage to containers.
- Create systemd service files to manage a container.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

```
[student@workstation ~]$ lab start containers-lifecycle
```

Instructions

Your organization requires that a web service that runs in a container should be a service that is managed by systemd. Configure a user account for systemd user services to start the container when the host machine starts.

- 1. Log in to the `servera` machine as the `contsvc` user. Use `redhat` as password.
 - 1.1. To manage the systemd user services with the `contsvc` account, you must log in directly as the `contsvc` user. You cannot use the `su` and `sudo` commands to create a session with the `contsvc` user.

```
[student@workstation ~]$ ssh contsvc@servera
...output omitted...
[contsvc@servera ~]$
```

- 2. Configure the `registry.lab.example.com` classroom registry in your home directory. Log in to the container registry with `admin` as the user and `redhat321` as the password.
 - 2.1. Create the `/home/contsvc/.config/containers` directory.

```
[contsvc@servera ~]$ mkdir -p /home/contsvc/.config/containers
```

- 2.2. Create the `/home/contsvc/.config/containers/registries.conf` file with the following contents:

```
unqualified-search-registries = ['registry.lab.example.com']

[[registry]]
location = "registry.lab.example.com"
insecure = true
blocked = false
```

- 2.3. Log in to the classroom registry.

```
[contsvc@servera ~]$ podman login registry.lab.example.com
Username: admin
Password: redhat321
Login Succeeded!
```

- ▶ 3. Use the `/home/contsvc/webcontent/html/` directory as persistent storage for the web server container. Create the `index.html` test page with the `Hello World` line inside the directory.

- 3.1. Create the `~/webcontent/html/` directory.

```
[contsvc@servera ~]$ mkdir -p ~/webcontent/html/
```

- 3.2. Create the `index.html` file and add the `Hello World` line.

```
[contsvc@servera ~]$ echo "Hello World" > ~/webcontent/html/index.html
```

- 3.3. Verify that the permission for others is set to `r-x` in the `webcontent/html` directory, and is set to `r--` in the `index.html` file. The container uses a non-privileged user that must be able to read the `index.html` file.

```
[contsvc@servera ~]$ ls -ld webcontent/html/
drwxr-xr-x. 2 contsvc contsvc 24 Jun  6 19:12 webcontent/html/
[contsvc@servera ~]$ ls -l webcontent/html/index.html
-rw-r--r--. 1 contsvc contsvc 12 Jun  6 19:12 webcontent/html/index.html
```

- ▶ 4. Create a container that runs an Apache HTTP server in the background. Use the `registry.lab.example.com/rhel9/httpd-24` image to run a container called `webapp` in detached mode. Redirect the `8090` port on the local host to the container `8080` port. Mount the `~/webcontent` directory from the host to the `/var/www` directory in the container.

- 4.1. Execute the `podman run` command to create the container. Use the `registry.lab.example.com/rhel9/httpd-24` image to run a container called `webapp` in detached mode. Use the `-p` option to map the `8090` port on `servera` to the `8080` port in the container. Use the `-v` option to mount the `~/webcontent` directory on `servera` to the `/var/www` directory in the container. Use the `Z` option to set the SELinux context to the mounted directory.

```
[contsvc@servera ~]$ podman run -d --name webapp -p 8090:8080 \
-v ~/webcontent:/var/www:Z registry.lab.example.com/rhel9/httpd-24
Trying to pull registry.lab.example.com/rhel9/httpd-24:latest...
...output omitted...
79ed4591cb59bdfba1d38badcf52289ee70213bafd7a1a2aefa4d2963cf88e29
```

4.2. Verify that the container is running. Use `podman ps` to list all the running containers.

```
[contsvc@servera ~]$ podman ps
```

CONTAINER ID	IMAGE	COMMAND
79ed4591cb59	registry.lab.example.com/rhel9/httpd-24:latest	/usr/bin/run-http...

```
3 minutes ago Up 3 minutes 0.0.0.0:8090->8080/tcp webapp
```

4.3. Use the `podman inspect` command to get the `Status` field, which indicates whether the container is running.

```
[contsvc@servera ~]$ podman inspect --format='{{.State.Status}}' webapp
running
```

4.4. Verify that the container is running by using the `Running` field.

```
[contsvc@servera ~]$ podman inspect --format='{{.State.Running}}' webapp
true
```

4.5. Verify that the web service is working on port 8090.

```
[contsvc@servera ~]$ curl http://localhost:8090
Hello World
```

► 5. Before creating a systemd unit file for the service container, test the functionality of the `webapp` container.

5.1. Stop the container. Use the container name to stop the container.

```
[contsvc@servera ~]$ podman stop webapp
webapp
```

5.2. Verify that the container is not running.

```
[contsvc@servera ~]$ podman inspect --format='{{.State.Status}}' webapp
exited
[contsvc@servera ~]$ podman inspect --format='{{.State.Running}}' webapp
false
```

5.3. Restart the container. Use the `podman restart` command to restart the container.

```
[student@workstation ~]$ podman restart webapp
webapp
```

5.4. Verify that the container is running again.

```
[contsvc@servera ~]$ podman ps
CONTAINER ID  IMAGE                                     COMMAND
CREATED      STATUS      PORTS      NAMES
79ed4591cb59  registry.lab.example.com/rhel9/httpd-24:latest /usr/bin/run-http...
20 minutes ago Up 29 seconds 0.0.0.0:8090->8080/tcp webapp
[contsvc@servera ~]$ curl http://localhost:8090
Hello World
```

- ▶ 6. Create a `systemd` service file to manage the `webapp` container with `systemctl` commands. Configure the `systemd` service so that when you start the service, the `systemd` daemon creates a container. After you finish the configuration, stop and then delete the `webapp` container. Remember that the `systemd` daemon expects that the container does not exist initially.

6.1. Create and change to the `~/.config/systemd/user/` directory.

```
[contsvc@servera ~]$ mkdir -p ~/.config/systemd/user/
[contsvc@servera ~]$ cd ~/.config/systemd/user
```

6.2. Create the unit file for the `webapp` container. Use the `--new` option so that `systemd` creates a container when starting the service, and deletes the container when stopping the service.

```
[contsvc@servera user]$ podman generate systemd --new --files --name webapp
/home/contsvc/.config/systemd/user/container-webapp.service
```

6.3. Stop and then delete the `webapp` container.

```
[contsvc@servera user]$ podman stop webapp
webapp
[contsvc@servera user]$ podman rm webapp
webapp
[contsvc@servera user]$ podman ps -a
CONTAINER ID  IMAGE                                     COMMAND      CREATED      STATUS      PORTS      NAMES
```

- ▶ 7. Reload the `systemd` daemon configuration, and then enable and start your new `container-webapp` user service. Verify the `systemd` service configuration, stop and start the service, and display the web server response and the container status.

7.1. Reload the configuration to recognize the new unit file.

```
[contsvc@servera user]$ systemctl --user daemon-reload
```

7.2. Enable and start the `container-webapp` service.

```
[contsvc@servera user]$ systemctl --user enable --now container-webapp
Created symlink /home/contsvc/.config/systemd/user/default.target.wants/container-webapp.service → /home/contsvc/.config/systemd/user/container-webapp.service.
```

7.3. Verify that the web server responds to requests.

```
[contsvc@servera user]$ curl http://localhost:8090
Hello World
```

7.4. Verify that the container is running.

```
[contsvc@servera user]$ podman ps
CONTAINER ID   IMAGE                                     COMMAND
CREATED        STATUS        PORTS        NAMES
f6b209f0c915   registry.lab.example.com/rhel9/httpd-24:latest /usr/bin/run-http...
About a minute ago Up About a minute 0.0.0.0:8090->8080/tcp webapp
```

Use the container ID information to confirm that the `systemd` daemon creates a container when you restart the service.

7.5. Stop the `container-webapp` service, and confirm that the container no longer exists. When you stop the service, the `systemd` daemon stops and then deletes the container.

```
[contsvc@servera user]$ systemctl --user stop container-webapp
[contsvc@servera user]$ podman ps --all
CONTAINER ID   IMAGE   COMMAND   CREATED   STATUS   PORTS   NAMES
```

7.6. Start the `container-webapp` service, and then confirm that the container is running.

The container ID is different, because the `systemd` daemon creates a container with the start instruction, and deletes the container with the stop instruction.

```
[contsvc@servera user]$ systemctl --user start container-webapp
[contsvc@servera user]$ podman ps
CONTAINER ID   IMAGE                                     COMMAND
CREATED        STATUS        PORTS        NAMES
f9488c4e4c9e   registry.lab.example.com/rhel9/httpd-24:latest /usr/bin/run-http...
5 seconds ago Up 6 seconds 0.0.0.0:8090->8080/tcp webapp
```

► 8. Ensure that the services for the `contsvc` user start at system boot. When done, restart the `servera` machine.

8.1. Run the `loginctl enable-linger` command.

```
[contsvc@servera user]$ loginctl enable-linger
```

8.2. Confirm that the `Linger` option is set for the `contsvc` user.

```
[contsvc@servera user]$ loginctl show-user contsvc
...output omitted...
Linger=yes
```

8.3. Switch to the `root` user, and then use the `systemctl reboot` command to restart the `servera` machine.

```
[contsvc@servera user]$ su -
Password: redhat
Last login: Wed Apr 17 07:44:06 EDT 2024 on tty2
[root@servera ~]# systemctl reboot
Connection to servera closed by remote host.
Connection to servera closed.
[student@workstation ~]$
```

- 9. When the `servera` machine is running again, log in to `servera` as the `contsvc` user. Verify that the `systemd` daemon started the `webapp` container, and that the web content is available.

9.1. Log in to the `servera` machine as the `contsvc` user.

```
[student@workstation ~]$ ssh contsvc@servera
```

9.2. Verify that the container is running.

```
[contsvc@servera ~]$ podman ps
```

CONTAINER ID	IMAGE	STATUS	PORTS	NAMES	COMMAND
694eb00c7583	registry.lab.example.com/rhel9/httpd-24:latest	Up 2 minutes ago	0.0.0.0:8090->8080/tcp	webapp	/usr/bin/run-http...

9.3. Access the web content.

```
[contsvc@servera ~]$ curl http://localhost:8090
Hello World
```

9.4. Return to the `workstation` machine as the `student` user.

```
[contsvc@servera ~]$ exit
logout
Connection to servera closed.
[student@workstation ~]$
```

Finish

On the `workstation` machine, use the `lab` command to complete this exercise. This is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish containers-lifecycle
```


► Lab

Run Containers

Use podman to pull a container image from a registry, and use the image to run a detached container.

Outcomes

- Create rootless detached containers.
- Configure a container image registry and create a container from an existing image.
- Configure port mapping and persistent storage.
- Configure a container as a `systemd` service and use the `systemctl` command to manage it.

Before You Begin

As the student user on the workstation machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start containers-review
```

Instructions

1. On the `serverb` machine, configure the `podsvc` user with `redhat` as the password. Configure the `registry.lab.example.com` registry as the remote registry. Use `admin` as the user and `redhat321` as the password to authenticate.
2. Configure the `registry.lab.example.com` classroom registry in your home directory. Log in to the container registry with `admin` as the user and `redhat321` as the password.
3. Use the `/home/student/webserver/html/` directory as persistent storage for the web server container. Create the `index.html` test page with the `Welcome to the webserver` container content.
4. Create a container that runs an Apache HTTP server in the background. Use the `registry.lab.example.com/rhel9/httpd-24` image to run a detached container called `webserver`. Redirect the 8080 port on the local host to the container 8080 port. Mount the `~/webserver` directory from the host to the `/var/www` directory in the container.
5. Create the `systemd` service file to manage the `webserver` container with `systemctl` commands. Configure the `systemd` service so that when you start the service, the `systemd` daemon creates a container. The `systemd` daemon expects that the container does not exist initially.
6. Reload the `systemd` daemon configuration, and then enable and start your new `container-webserver` user service. Verify that the `webserver` container is started and running.

7. Ensure that the services for the `podsvc` user start at system boot.
8. Verify that the web service is working on the 8080 port and that the content is accessible.
9. Return to the `workstation` machine as the `student` user.

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade containers-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish containers-review
```

► Solution

Run Containers

Use podman to pull a container image from a registry, and use the image to run a detached container.

Outcomes

- Create rootless detached containers.
- Configure a container image registry and create a container from an existing image.
- Configure port mapping and persistent storage.
- Configure a container as a `systemd` service and use the `systemctl` command to manage it.

Before You Begin

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start containers-review
```

Instructions

1. On the `serverb` machine, configure the `podsvc` user with `redhat` as the password. Configure the `registry.lab.example.com` registry as the remote registry. Use `admin` as the user and `redhat321` as the password to authenticate.

- 1.1. Log in to the `serverb` machine as the `student` user.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
[student@serverb ~]$
```

- 1.2. Create the `podsvc` user and set `redhat` as the password for the user. Use `student` as the password for creating the user by using the `sudo` command.

```
[student@serverb ~]$ sudo useradd podsvc
[sudo] password for student: student
[student@serverb ~]$ sudo passwd podsvc
Changing password for user podsvc.
New password: redhat
BAD PASSWORD: The password is shorter than 8 characters
Retype new password: redhat
passwd: all authentication tokens updated successfully.
```

- 1.3. Return to the workstation machine as the student user.

```
[student@serverb ~]$ exit
logout
Connection to serverb closed.
[student@workstation ~]$
```

- 1.4. Log in to the serverb machine as the podsvc user. Use redhat as the password.

```
[student@workstation ~]$ ssh podsvc@serverb
...output omitted...
[podsvc@serverb ~]$
```

2. Configure the registry.lab.example.com classroom registry in your home directory. Log in to the container registry with admin as the user and redhat321 as the password.

- 2.1. Create the /home/podsvc/.config/containers directory.

```
[podsvc@serverb ~]$ mkdir -p /home/podsvc/.config/containers
```

- 2.2. Create the /home/podsvc/.config/containers/registries.conf file with the following contents:

```
unqualified-search-registries = ['registry.lab.example.com']

[[registry]]
location = "registry.lab.example.com"
insecure = true
blocked = false
```

- 2.3. Log in to the classroom registry.

```
[podsvc@serverb ~]$ podman login registry.lab.example.com
Username: admin
Password: redhat321
Login Succeeded!
```

3. Use the /home/student/webserver/html/ directory as persistent storage for the web server container. Create the index.html test page with the Welcome to the webserver container content.

- 3.1. Create the ~/webserver/html/ directory.

```
[podsvc@serverb ~]$ mkdir -p ~/webserver/html/
```

- 3.2. Create the index.html file and add the Welcome to the webserver container content.

```
[podsvc@serverb ~]$ echo "Welcome to the webserver container" > \
~/webserver/html/index.html
```

- 3.3. Verify that the permission for others is set to `r-x` in the `webserver/html` directory, and is set to `r--` in the `index.html` file. The container uses a non-privileged user that must be able to read the `index.html` file.

```
[podsvc@serverb ~]$ ls -ld ~/webserver/html/
drwxr-xr-x. 2 podsvc podsvc 24 Jul 10 05:42 /home/podsvc/webserver/html/
[podsvc@serverb ~]$ ls -l ~/webserver/html/index.html
-rw-r--r--. 1 podsvc podsvc 21 Jul 10 05:42 /home/podsvc/webserver/html/index.html
```

4. Create a container that runs an Apache HTTP server in the background. Use the `registry.lab.example.com/rhel9/httpd-24` image to run a detached container called `webserver`. Redirect the `8080` port on the local host to the container `8080` port. Mount the `~/webserver` directory from the host to the `/var/www` directory in the container.
 - 4.1. Execute the `podman run` command to create the detached container. Use the `registry.lab.example.com/rhel9/httpd-24` image to run a detached container called `webserver`. Use the `-p` option to map the `8080` port on the `serverb` machine to the `8080` port in the container. Use the `-v` option to mount the `~/webserver` directory on the `serverb` machine to the `/var/www` directory in the container. Use the `Z` option to set the SELinux context to the mounted directory.

```
[podsvc@serverb ~]$ podman run -d --name webserver -p 8080:8080 \
-v ~/webserver:/var/www:Z registry.lab.example.com/rhel9/httpd-24
...output omitted...
d970ff062f002a45702b96c0a51d632d93d78ccf63a3af1a01abf70bc4c46616
```

- 4.2. Verify that the container is running.

```
[podsvc@serverb ~]$ podman ps
```

CONTAINER ID	IMAGE	COMMAND
d970ff062f00	registry.lab.example.com/rhel9/httpd-24:latest	/usr/bin/run-http...

```
About a minute ago Up About a minute 0.0.0.0:8080->8080/tcp webserver
```

- 4.3. Verify that the web service is working on the `8080` port.

```
[podsvc@serverb ~]$ curl http://localhost:8080
Welcome to the webserver container
```

5. Create the `systemd` service file to manage the `webserver` container with `systemctl` commands. Configure the `systemd` service so that when you start the service, the `systemd` daemon creates a container. The `systemd` daemon expects that the container does not exist initially.

- 5.1. Create and change to the `~/config/systemd/user/` directory.

```
[podsvc@serverb ~]$ mkdir -p ~/.config/systemd/user/
[podsvc@serverb ~]$ cd ~/.config/systemd/user
```

- 5.2. Create the unit file for the `webserver` container. Use the `--new` option so that the `systemd` service creates a container when starting the service, and deletes the container when stopping the service.

```
[podsvc@serverb user]$ podman generate systemd --new --files --name webserver
/home/podsvc/.config/systemd/user/container-webserver.service
```

5.3. Stop and then delete the `webserver` container.

```
[podsvc@serverb user]$ podman stop webserver
webserver
[podsvc@serverb user]$ podman rm webserver
webserver
[podsvc@serverb user]$ podman ps -a
```

CONTAINER ID	IMAGE	COMMAND	CREATED	STATUS	PORTS	NAMES

6. Reload the `systemd` daemon configuration, and then enable and start your new `container-webserver` user service. Verify that the `webserver` container is started and running.

6.1. Reload the configuration to recognize the new unit file.

```
[podsvc@serverb user]$ systemctl --user daemon-reload
```

6.2. Enable and start the `container-webserver` service.

```
[podsvc@serverb user]$ systemctl --user enable --now container-webserver
Created symlink /home/podsvc/.config/systemd/user/default.target.wants/container-
webserver.service → /home/podsvc/.config/systemd/user/container-webserver.service.
```

6.3. Verify that the container is running.

```
[podsvc@serverb user]$ podman ps
```

CONTAINER ID	IMAGE	COMMAND
4425565b3192	registry.lab.example.com/rhel9/httpd-24:latest	/usr/bin/run-http...

```
23 seconds ago Up 23 seconds 0.0.0.0:8080->8080/tcp webserver
```

7. Ensure that the services for the `podsvc` user start at system boot.

7.1. Run the `loginctl enable-linger` command.

```
[podsvc@serverb user]$ loginctl enable-linger
```

7.2. Confirm that the `Linger` option is set for the `podsvc` user.

```
[podsvc@serverb user]$ loginctl show-user podsvc
...output omitted...
Linger=yes
```

8. Verify that the web service is working on the 8080 port and that the content is accessible.

```
[podsvc@serverb user]$ curl http://localhost:8080
Welcome to the webserver container
```

9. Return to the `workstation` machine as the `student` user.

```
[podsvc@serverb ~]$ exit
logout
Connection to serverb closed.
[student@workstation ~]$
```

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade containers-review
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish containers-review
```

Summary

- A container is an encapsulated process that includes the required runtime dependencies for the program to run.
- Registries are used to store container images to later share them in a controlled manner.
- Containers are ephemeral by default, which means that the container engine removes the writable layer when you remove the container.
- Hypervisors are applications that provide the virtualization functionality for virtual machines. The container equivalent of a hypervisor is a container engine, such as Podman.
- Podman, which Red Hat Enterprise Linux provides, is a container engine that directly runs and manages container instances and container images on a single host.
- With Podman, you can find, run, build, or deploy OCI (Open Container Initiative) containers and container images. By default, Podman is daemonless.
- As a regular user, you can create systemd unit files to manage your rootless containers.

Chapter 14

Comprehensive Review

Goal

Review tasks from *Red Hat System Administration II*.

Sections

- Comprehensive Review

Lab

- Fix Boot Issues and Maintain Servers
- Configure and Manage File Systems and Storage
- Configure and Manage Server Security
- Run Containers

Comprehensive Review

Objectives

After completing this section, you should have reviewed and refreshed the knowledge and skills that you learned in *Red Hat System Administration II*.

Reviewing Red Hat System Administration II

Before beginning the comprehensive review for this course, you should be comfortable with the topics covered in each chapter. Do not hesitate to ask the instructor for extra guidance or clarification on these topics.

Chapter 1, Improve Command-line Productivity

Run commands more efficiently by using advanced features of the Bash shell, shell scripts, and various Red Hat Enterprise Linux utilities.

- Run commands more efficiently by using advanced features of the Bash shell, shell scripts, and various Red Hat Enterprise Linux utilities.
- Run repetitive tasks with `for` loops, evaluate exit codes from commands and scripts, run tests with operators, and create conditional structures with `if` statements.
- Create regular expressions to match data, apply regular expressions to text files with the `grep` command, and use `grep` to search files and data from piped commands.

Chapter 2, Schedule Future Tasks

Schedule tasks to execute at a specific time and date.

- Set up a command to run once at a future time.
- Schedule commands to run on a repeating schedule with a user's `crontab` file.
- Schedule commands to run on a repeating schedule with the system `crontab` file and directories.
- Enable and disable `systemd` timers, and configure a timer that manages temporary files.

Chapter 3, Analyze and Store Logs

Locate and accurately interpret system event logs for troubleshooting purposes.

- Describe the basic Red Hat Enterprise Linux logging architecture to record events.
- Interpret events in the relevant `syslog` files to troubleshoot problems or to review system status.
- Find and interpret entries in the system journal to troubleshoot problems or review system status.
- Configure the system journal to preserve the record of events when a server is rebooted.

- Maintain accurate time synchronization with Network Time Protocol (NTP) and configure the time zone to ensure correct time stamps for events that are recorded by the system journal and logs.

Chapter 4, Archive and Transfer Files

Archive and copy files from one system to another.

- Archive files and directories into a compressed file with `tar`, and extract the contents of an existing `tar` archive.
- Transfer files to or from a remote system securely with SSH.
- Efficiently and securely synchronize the contents of a local file or directory with a remote server copy.

Chapter 5, Tune System Performance

Improve system performance by setting tuning parameters and adjusting the scheduling priority of processes.

- Optimize system performance by selecting a tuning profile that the `tuned` daemon manages.
- Prioritize or deprioritize specific processes, with the `nice` and `renice` commands.

Chapter 6, Manage SELinux Security

Protect and manage server security by using SELinux.

- Explain how SELinux protects resources, change the current SELinux mode of a system, and set the default SELinux mode of a system.
- Manage the SELinux policy rules that determine the default context for files and directories with the `semanage fcontext` command, and apply the context defined by the SELinux policy to files and directories with the `restorecon` command.
- Activate and deactivate SELinux policy rules with the `setsebool` command, manage the persistent value of SELinux Booleans with the `semanage boolean -l` command, and consult `man` pages that end with `_selinux` to find useful information about SELinux Booleans.
- Use SELinux log analysis tools and display useful information during SELinux troubleshooting with the `sealert` command.

Chapter 7, Manage Basic Storage

Create and manage storage devices, partitions, file systems, and swap spaces from the command line.

- Create storage partitions, format them with file systems, and mount them for use.
- Create and manage swap spaces to supplement physical memory.

Chapter 8, Manage Storage Stack

Create and manage logical volumes that contain file systems or swap spaces from the command line.

- Describe logical volume manager components and concepts, and implement LVM storage and display LVM component information.

- Analyze the multiple storage components that make up the layers of the storage stack.

Chapter 9, Access Network-Attached Storage

Access network-attached storage with the NFS protocol.

- Identify NFS export information, create a directory to use as a mount point, mount an NFS export with the `mount` command or by configuring the `/etc/fstab` file, and unmount an NFS export with the `umount` command.
- Describe the benefits of using the automounter, and automount NFS exports by using direct and indirect maps.

Chapter 10, Control the Boot Process

Manage the boot process to control offered services and to troubleshoot and repair problems.

- Describe the Red Hat Enterprise Linux boot process, set the default target when booting, and boot a system to a non-default target.
- Log in to a system and change the root password when the current root password is lost.
- Manually repair file-system configuration or corruption issues that stop the boot process.

Chapter 11, Manage Network Security

Control network connections to services with the system firewall and SELinux rules.

- Accept or reject network connections to system services with `firewalld` rules.
- Verify that network ports have the correct SELinux type for services to bind to them.

Chapter 12, Install Red Hat Enterprise Linux

Install Red Hat Enterprise Linux on servers and virtual machines.

- Install Red Hat Enterprise Linux on a server.
- Explain Kickstart concepts and architecture, create a Kickstart file with the Kickstart Generator website, modify an existing Kickstart file with a text editor and check its syntax with `ksvalidator`, publish a Kickstart file to the installer, and install Kickstart on the network.
- Install a virtual machine on your Red Hat Enterprise Linux server with the web console.

Chapter 13, Run Containers

Obtain, run, and manage simple lightweight services as containers on a single Red Hat Enterprise Linux server.

- Explain container concepts and the core technologies for building, storing, and running containers.
- Pull container images from a registry and run containers by using Podman.
- Authenticate to and explore the container images stored in remote container registries.
- Manage the lifecycle of a container from creation to deletion.

► Lab

Fix Boot Issues and Maintain Servers



Note

If you plan to take the RHCSA exam, then use the following approach to maximize the benefit of this Comprehensive Review: attempt each lab without viewing the solution buttons or referring to the course content. Use the grading scripts to gauge your progress as you complete each lab.

Troubleshoot and repair boot problems and update the system default target. You also schedule tasks to run on a repeating schedule as a normal user.

Outcomes

- Diagnose issues and recover the system from emergency mode.
- Change the default target from `graphical.target` to `multi-user.target`.
- Schedule recurring jobs to run as a normal user.

Before You Begin

If you did not reset your `workstation` and `serverb` machines at the end of the last chapter, then save any work that you want to keep from earlier exercises on those machines, and reset them now.

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start rhcsa-compreview1
```

Specifications

- On `workstation`, run the `/tmp/rhcsa-break1` script. This script causes an issue with the boot process on `serverb` and then reboots the machine. Troubleshoot the cause and repair the boot issue. When prompted, use `redhat` as the password of the `root` user.
- On `workstation`, run the `/tmp/rhcsa-break2` script. This script causes the default target to switch from the `multi-user` target to the `graphical` target on the `serverb` machine and then reboots the machine. On `serverb`, reset the default target to use the `multi-user` target. The default target settings must persist after reboot without manual intervention. As the `student` user, use the `sudo` command for performing privileged commands. Use `student` as the password, when required.
- On `serverb`, schedule a recurring job as the `student` user that executes the `/home/student/backup-home.sh` script hourly between 7 PM and 9 PM every day except on

Saturday and Sunday. Download the backup script from `http://materials.example.com/labs/backup-home.sh`. The `backup-home.sh` script backs up the `/home/student` directory from `serverb` to `servera` in the `/home/student/serverb-backup` directory. Use the `backup-home.sh` script to schedule the recurring job as the `student` user. Run the command as an executable.

- Reboot the `serverb` machine and wait for the boot to complete before grading.

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade rhcsa-compreview1
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish rhcsa-compreview1
```

► Solution

Fix Boot Issues and Maintain Servers



Note

If you plan to take the RHCSA exam, then use the following approach to maximize the benefit of this Comprehensive Review: attempt each lab without viewing the solution buttons or referring to the course content. Use the grading scripts to gauge your progress as you complete each lab.

Troubleshoot and repair boot problems and update the system default target. You also schedule tasks to run on a repeating schedule as a normal user.

Outcomes

- Diagnose issues and recover the system from emergency mode.
- Change the default target from `graphical.target` to `multi-user.target`.
- Schedule recurring jobs to run as a normal user.

Before You Begin

If you did not reset your `workstation` and `serverb` machines at the end of the last chapter, then save any work that you want to keep from earlier exercises on those machines, and reset them now.

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start rhcsa-compreview1
```

1. On `workstation`, run the `/tmp/rhcsa-break1` script.

```
[student@workstation ~]$ sh /tmp/rhcsa-break1
```

2. After the `serverb` machine boots, access the console and notice that the boot process stopped early. Consider a possible cause for this behavior.
 - 2.1. Locate the icon for the `serverb` console, as appropriate for your classroom environment. Open the console and inspect the error. It might take a few seconds for the error to appear.
 - 2.2. Press `Ctrl+Alt+Del` to reboot the `serverb` machine. When the boot-loader menu appears, press any key except `Enter` to interrupt the countdown.

- 2.3. Edit the default boot-loader entry, in memory, to log in to the emergency mode. Press **e** to edit the current entry.
- 2.4. Use the cursor keys to navigate to the line that starts with **linux**. Append **systemd.unit=emergency.target**.
- 2.5. Press **Ctrl+x** to boot with the modified configuration.
- 2.6. Log in to emergency mode. Use **redhat** as the **root** user's password.

```
Give root password for maintenance
(or press Control-D to continue): redhat
[root@serverb ~]#
```

3. Remount the **/** file system with read and write capabilities. Use the **mount -a** command to try to mount all the other file systems.
 - 3.1. Remount the **/** file system with read and write capabilities to edit the file system.

```
[root@serverb ~]# mount -o remount,rw /
```

- 3.2. Try to mount all the other file systems. Notice that one of the file systems does not mount.

```
[root@serverb ~]# mount -a
...output omitted...
mount: /FakeMount: can't find UUID=fake.
```

- 3.3. Edit the **/etc/fstab** file to fix the issue. Remove or comment out the incorrect line.

```
[root@serverb ~]# vi /etc/fstab
...output omitted...
#UUID=fake      /FakeMount    xfs     defaults    0 0
```

- 3.4. Update the **systemd** daemon for the system to register the new **/etc/fstab** file configuration.

```
[root@serverb ~]# systemctl daemon-reload
[ 206.828912] systemd[1]: Reloading.
```

- 3.5. Verify that **/etc/fstab** file is now correct by attempting to mount all entries.

```
[root@serverb ~]# mount -a
```

- 3.6. Reboot **serverb** and wait for the boot to complete. The system should now boot without errors.

```
[root@serverb ~]# systemctl reboot
```

4. On workstation, run the **/tmp/rhcsa-break2** script. Wait for the **serverb** machine to reboot before proceeding.

```
[student@workstation ~]$ sh /tmp/rhcsa-break2
```

5. On `serverb`, set the `multi-user` target as the current and default target.

- 5.1. Log in to `serverb` as the `student` user.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
[student@serverb ~]$
```

- 5.2. Determine the default target.

```
[student@serverb ~]$ systemctl get-default
graphical.target
```

- 5.3. Switch to the `multi-user` target.

```
[student@serverb ~]$ sudo systemctl isolate multi-user.target
[sudo] password for student: student
```

- 5.4. Set the `multi-user` target as the default target.

```
[student@serverb ~]$ sudo systemctl set-default multi-user.target
Removed /etc/systemd/system/default.target.
Created symlink /etc/systemd/system/default.target -> /usr/lib/systemd/system/multi-user.target.
```

- 5.5. Reboot `serverb` and verify that the `multi-user` target is set as the default target.

```
[student@serverb ~]$ sudo systemctl reboot
Connection to serverb closed by remote host.
Connection to serverb closed.
[student@workstation ~]$
```

- 5.6. After the system reboots, open an SSH session to `serverb` as the `student` user. Verify that the `multi-user` target is set as the default target.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
[student@serverb ~]$ systemctl get-default
multi-user.target
```

6. On `serverb`, schedule a recurring job as the `student` user that executes the `/home/student/backup-home.sh` script hourly between 7 PM and 9 PM on all days except Saturday and Sunday. Use the `backup-home.sh` script to schedule the recurring job. Download the backup script from `http://materials.example.com/labs/backup-home.sh`. Run the command as an executable.
 - 6.1. On `serverb`, download the backup script from `http://materials.example.com/labs/backup-home.sh`. Use `chmod` to make the backup script executable.

```
[student@serverb ~]$ wget http://materials.example.com/labs/backup-home.sh
...output omitted...
[student@serverb ~]$ chmod +x backup-home.sh
```

6.2. Open the crontab file with the default text editor.

```
[student@serverb ~]$ crontab -e
```

6.3. Edit the file to add the following line:

```
0 19-21 * * Mon-Fri /home/student/backup-home.sh
```

Save the changes and exit the editor.

6.4. Use the `crontab -l` command to list the scheduled recurring jobs.

```
[student@serverb ~]$ crontab -l
0 19-21 * * Mon-Fri /home/student/backup-home.sh
```

7. Reboot `serverb` and wait for the boot to complete before grading.

```
[student@serverb ~]$ sudo systemctl reboot
[sudo] password for student: student
Connection to serverb closed by remote host.
Connection to serverb closed.
[student@workstation ~]$
```

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade rhcsa-compreview1
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish rhcsa-compreview1
```

► Lab

Configure and Manage File Systems and Storage



Note

If you plan to take the RHCSA exam, then use the following approach to maximize the benefit of this Comprehensive Review: attempt each lab without viewing the solution buttons or referring to the course content. Use the grading scripts to gauge your progress as you complete each lab.

Create a logical volume, mount a network file system, and create a swap partition that is automatically activated at boot. You also configure directories to store temporary files.

Outcomes

- Create a logical volume.
- Mount a network file system.
- Create a swap partition that is automatically activated at boot.
- Configure a directory to store temporary files.

Before You Begin

If you did not reset your **workstation** and **server** machines at the end of the last chapter, then save any work that you want to keep from earlier exercises on those machines, and reset them now.

As the **student** user on the **workstation** machine, use the **lab** command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start rhcsa-compreview2
```

Specifications

- On **serverb**, configure a new 1 GiB **vo1_home** logical volume in a new 2 GiB **extra_storage** volume group. Use the unpartitioned **/dev/vdb** disk to create the partition.
- Format the **vo1_home** logical volume with the **XFS** file-system type, and persistently mount it on the **/user-homes** directory.
- On **serverb**, persistently mount the **/share** network file system that **servera** exports on the **/local-share** directory. The **servera** machine exports the **servera.lab.example.com:/share** path.

- On `serverb`, create a 512 MiB swap partition on the `/dev/vdc` disk. Persistently mount the swap partition.
- Create the `production` user group. Create the `production1`, `production2`, `production3`, and `production4` users with the `production` group as their supplementary group.
- On `serverb`, configure the `/run/volatile` directory to store temporary files. If the files in this directory are not accessed for more than 30 seconds, then the system automatically deletes them. Set `0700` as the octal permissions for the directory. Use the `/etc/tmpfiles.d/volatile.conf` file to configure the time-based deletion of the files in the `/run/volatile` directory.

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade rhcsa-compreview2
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish rhcsa-compreview2
```

► Solution

Configure and Manage File Systems and Storage



Note

If you plan to take the RHCSA exam, then use the following approach to maximize the benefit of this Comprehensive Review: attempt each lab without viewing the solution buttons or referring to the course content. Use the grading scripts to gauge your progress as you complete each lab.

Create a logical volume, mount a network file system, and create a swap partition that is automatically activated at boot. You also configure directories to store temporary files.

Outcomes

- Create a logical volume.
- Mount a network file system.
- Create a swap partition that is automatically activated at boot.
- Configure a directory to store temporary files.

Before You Begin

If you did not reset your `workstation` and `server` machines at the end of the last chapter, then save any work that you want to keep from earlier exercises on those machines, and reset them now.

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start rhcsa-compreview2
```

1. On `serverb`, configure a new 1 GiB `vol_home` logical volume in a new 2 GiB `extra_storage` volume group. Use the unpartitioned `/dev/vdb` disk to create the partition.
 - 1.1. Log in to `serverb` as the `student` user and switch to the `root` user.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
[student@serverb ~]$ sudo -i
[sudo] password for student: student
[root@serverb ~]#
```

- 1.2. Create a 2 GiB partition on the /dev/vdb disk.

```
[root@serverb ~]# parted /dev/vdb mklabel msdos
...output omitted...
[root@serverb ~]# parted /dev/vdb mkpart primary 1MiB 2GiB
...output omitted...
[root@serverb ~]# parted /dev/vdb set 1 lvm on
...output omitted...
```

- 1.3. Declare the /dev/vdb1 block device as a physical volume.

```
[root@serverb ~]# pvcreate /dev/vdb1
...output omitted...
```

- 1.4. Create the extra_storage volume group with the /dev/vdb1 partition.

```
[root@serverb ~]# vgcreate extra_storage /dev/vdb1
...output omitted...
```

- 1.5. Create the 1 GiB vol_home logical volume.

```
[root@serverb ~]# lvcreate -L 1GiB -n vol_home extra_storage
...output omitted...
```

2. Format the vol_home logical volume with the XFS file-system type, and persistently mount it on the /user-homes directory.

- 2.1. Create the /user-homes directory.

```
[root@serverb ~]# mkdir /user-homes
```

- 2.2. Format the /dev/extra_storage/vol_home partition with the XFS file-system type.

```
[root@serverb ~]# mkfs -t xfs /dev/extra_storage/vol_home
...output omitted...
```

- 2.3. Persistently mount the /dev/extra_storage/vol_home partition on the /user-homes directory. Use the partition's UUID for the /etc/fstab file entry.

```
[root@serverb ~]# lsblk -o UUID /dev/extra_storage/vol_home
UUID
988cf149-0667-4733-abca-f80c6ec50ab6
[root@serverb ~]# echo "UUID=988c...0ab6 /user-homes xfs defaults 0 0" \
>> /etc/fstab
[root@serverb ~]# mount /user-homes
```

3. On serverb, persistently mount the /share network file system that servera exports on the /local-share directory. The servera machine exports the servera.lab.example.com:/share path.

- 3.1. Create the /local-share directory.

```
[root@serverb ~]# mkdir /local-share
```

- 3.2. Append the appropriate entry to the /etc/fstab file to persistently mount the servera.lab.example.com:/share network file system.

```
[root@serverb ~]# echo "servera.lab.example.com:/share /local-share \
nfs rw, sync 0 0" >> /etc/fstab
```

- 3.3. Mount the network file system on the /local-share directory.

```
[root@serverb ~]# mount /local-share
```

4. On serverb, create a 512 MiB swap partition on the /dev/vdc disk. Activate and persistently mount the swap partition.

- 4.1. Create a 512 MiB partition on the /dev/vdc disk.

```
[root@serverb ~]# parted /dev/vdc mklabel msdos
...output omitted...
[root@serverb ~]# parted /dev/vdc mkpart primary linux-swaps 1MiB 513MiB
...output omitted...
```

- 4.2. Create the swap space on the /dev/vdc1 partition.

```
[root@serverb ~]# mkswap /dev/vdc1
...output omitted...
```

- 4.3. Create an entry in the /etc/fstab file to persistently mount the swap space. Use the partition's UUID to create the /etc/fstab file entry. Activate the swap space.

```
[root@serverb ~]# lsblk -o UUID /dev/vdc1
UUID
cc18ccb6-bd29-48a5-8554-546bf3471b69
[root@serverb ~]# echo "UUID=cc18...1b69 swap swap defaults 0 0" >> /etc/fstab
[root@serverb ~]# swapon -a
```


5. Create the `production` user group. Then, create the `production1`, `production2`, `production3`, and `production4` users with the `production` group as their supplementary group.

```
[root@serverb ~]# groupadd production
[root@serverb ~]# for i in 1 2 3 4; do useradd -G production production$i; done
```

6. On `serverb`, configure the `/run/volatile` directory to store temporary files. If the files in this directory are not accessed for more than 30 seconds, then the system automatically deletes them. Set `0700` as the octal permissions for the directory. Use the `/etc/tmpfiles.d/volatile.conf` file to configure the time-based deletion of the files in the `/run/volatile` directory.

6.1. Create the `/etc/tmpfiles.d/volatile.conf` file with the following content:

```
d /run/volatile 0700 root root 30s
```

6.2. Use the `systemd-tmpfiles --create` command to create the `/run/volatile` directory if it does not exist.

```
[root@serverb ~]# systemd-tmpfiles --create /etc/tmpfiles.d/volatile.conf
```

6.3. Return to the `workstation` machine as the `student` user.

```
[root@serverb ~]# exit
logout
[student@serverb ~]$ exit
logout
Connection to serverb closed.
```

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade rhcsa-compreview2
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish rhcsa-compreview2
```

► Lab

Configure and Manage Server Security



Note

If you plan to take the RHCSA exam, then use the following approach to maximize the benefit of this Comprehensive Review: attempt each lab without viewing the solution buttons or referring to the course content. Use the grading scripts to gauge your progress as you complete each lab.

Configure SSH key-based authentication, change firewall settings, adjust the SELinux mode and an SELinux Boolean, and troubleshoot SELinux issues.

Outcomes

- Configure SSH key-based authentication.
- Configure firewall settings.
- Adjust the SELinux mode and SELinux Booleans.
- Troubleshoot SELinux issues.

Before You Begin

If you did not reset your `workstation` and `server` machines at the end of the last chapter, then save any work that you want to keep from earlier exercises on those machines, and reset them now.

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start rhcsa-compreview3
```

Specifications

- On `serverb`, generate an SSH key pair for the `student` user. Do not protect the private key with a passphrase.
- Configure the `student` user on `servera` to accept login authentication with the SSH key pair that you generated on the `serverb` machine. The `student` user on `serverb` must be able to log in to `servera` via SSH without entering a password.
- On `servera`, check the `/user-homes/production5` directory permissions. Then, configure SELinux to run in the `permissive` mode by default.
- On `serverb`, verify that the `/localhome` directory does not exist. Then, configure the `production5` user's home directory to mount the `/user-homes/production5` network

file system. The `servera.lab.example.com` machine exports the file system as the `servera.lab.example.com:/user-homes/production5` NFS share. Use the `autofs` service to mount the network share. Verify that the `autofs` service creates the `/localhome/production5` directory with the same permissions as on `servera`.

- On `serverb`, adjust the appropriate SELinux Boolean so that the `production5` user may use the NFS-mounted home directory after authenticating with an SSH key. If required, use `redhat` as the password of the `production5` user.
- On `serverb`, adjust the firewall settings to block all connection requests from the `servera` machine. Use the `servera` IPv4 address (`172.25.250.10`) to configure the firewall rule.
- On `serverb`, investigate and fix the issue with the failing Apache web service, which listens on port `30080/TCP` for connections. Adjust the firewall settings appropriately so that the port `30080/TCP` is open for incoming connections.

Evaluation

As the student user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade rhcsa-compreview3
```

Finish

On the `workstation` machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish rhcsa-compreview3
```

► Solution

Configure and Manage Server Security



Note

If you plan to take the RHCSA exam, then use the following approach to maximize the benefit of this Comprehensive Review: attempt each lab without viewing the solution buttons or referring to the course content. Use the grading scripts to gauge your progress as you complete each lab.

Configure SSH key-based authentication, change firewall settings, adjust the SELinux mode and an SELinux Boolean, and troubleshoot SELinux issues.

Outcomes

- Configure SSH key-based authentication.
- Configure firewall settings.
- Adjust the SELinux mode and SELinux Booleans.
- Troubleshoot SELinux issues.

Before You Begin

If you did not reset your `workstation` and `server` machines at the end of the last chapter, then save any work that you want to keep from earlier exercises on those machines, and reset them now.

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start rhcsa-compreview3
```

1. On `serverb`, generate an SSH key pair for the `student` user. Do not protect the private key with a passphrase.

- 1.1. Log in to `serverb` as the `student` user.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
```

- 1.2. Use the `ssh-keygen` command to generate an SSH key pair. Do not protect the private key with a passphrase.

```
[student@serverb ~]$ ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/home/student/.ssh/id_rsa): Enter
Created directory '/home/student/.ssh'.
Enter passphrase (empty for no passphrase): Enter
Enter same passphrase again: Enter
Your identification has been saved in /home/student/.ssh/id_rsa.
Your public key has been saved in /home/student/.ssh/id_rsa.pub.
The key fingerprint is:
SHA256:+ijpGqjEQSGBR80RNchiRTHw/URQksVdHjsHqVBXeYI student@serverb.lab.example.com
The key's randomart image is:
+---[RSA 3072]-----+
|+BBX+O*+o..=+.. |
|+.0.0000 .oE+o . |
|. + . . . . + .o |
|. . . . . o . o |
|. . . . . S |
|... . |
|.o. .. |
|o .o o |
|. .o... . |
+----[SHA256]-----+
```

2. Configure the `student` user on `servera` to accept login authentication with the SSH key pair that you generated on the `serverb` machine. The `student` user on `serverb` must be able to log in to `servera` via SSH without entering a password.

- 2.1. Send the public key of the newly generated SSH key pair to the `student` user on the `servera` machine.

```
[student@serverb ~]$ ssh-copy-id student@servera
/usr/bin/ssh-copy-id: INFO: Source of key(s) to be installed: "/home/student/.ssh/id_rsa.pub"
The authenticity of host 'servera (172.25.250.10)' can't be established.
ED25519 key fingerprint is SHA256:shYfoFG0Nnv42pv7j+HG+FISmCam4Bh5jfjwwSMJbrw.
This key is not known by any other names
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter
out any that are already installed
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompted
now it is to install the new keys
student@servera's password: student

Number of key(s) added: 1

Now try logging in to the machine, with: "ssh 'student@servera'"
and check to make sure that only the key(s) you wanted were added.
```

- 2.2. Verify that the `student` user can log in to `servera` from `serverb` without entering a password. Do not close the connection.

```
[student@serverb ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
```

3. On `servera`, verify the `/user-homes/production5` directory permissions. Then, configure SELinux to run in the permissive mode by default.

- 3.1. Verify the `/user-homes/production5` directory permissions.

```
[student@servera ~]$ ls -ld /user-homes/production5
drwx-----. 2 production5 production5 62 May  6 05:27 /user-homes/production5
```

- 3.2. Edit the `/etc/sysconfig/selinux` file to set the SELINUX parameter to the permissive value.

```
[student@servera ~]$ sudo vi /etc/sysconfig/selinux
...output omitted...
#SELINUX=enforcing
SELINUX=permissive
...output omitted...
```

- 3.3. Reboot the system.

```
[student@servera ~]$ sudo systemctl reboot
Connection to servera closed by remote host.
Connection to servera closed.
[student@serverb ~]$
```

4. On `serverb`, verify that the `/localhome` directory does not exist. Then, configure the `production5` user's home directory to mount the `/user-homes/production5` network file system. The `servera.lab.example.com` machine exports the file system as the `servera.lab.example.com:/user-homes/production5` NFS share. Use the `autofs` service to mount the network share. Verify that the `autofs` service creates the `/localhome/production5` directory with the same permissions as on `servera`.

- 4.1. Verify that the `/localhome` directory does not exist.

```
[student@serverb ~]$ ls -ld /localhome
ls: cannot access '/localhome': No such file or directory
```

- 4.2. On `serverb`, switch to the root user.

```
[student@serverb ~]$ sudo -i
[sudo] password for student: student
[root@serverb ~]#
```

- 4.3. Install the `autofs` package.

```
[root@serverb ~]# dnf install autofs
...output omitted...
Is this ok [y/N]: y
...output omitted...
Installed:
  autofs-1:5.1.7-27.el9.x86_64      libsss_autofs-2.6.2-2.el9.x86_64

Complete!
```

- 4.4. Create the `/etc/auto.master.d/production5.autofs` map file with the following content:

```
/- /etc/auto.production5
```

- 4.5. Determine the `production5` user's home directory.

```
[root@serverb ~]# getent passwd production5
production5:x:5001:5001::/localhome/production5:/bin/bash
```

- 4.6. Create the `/etc/auto.production5` file with the following content:

```
/localhome/production5 -rw servera.lab.example.com:/user-homes/production5
```

- 4.7. Restart the `autofs` service.

```
[root@serverb ~]# systemctl restart autofs
```

- 4.8. Verify that the `autofs` service creates the `/localhome/production5` directory on `serverb` with the same permissions as the `/user-homes/production5` directory on `servera`.

```
[root@serverb ~]# ls -ld /localhome/production5
drwx-----. 2 production5 production5 62 May  6 05:52 /localhome/production5
```

Note

The `autofs` service shows the correct permissions when you try to access the `/localhome/production5` shared directory. Before you access this directory, the file permissions incorrectly show that the directory has `root` as the user owner and the group owner. After you try to access the `/localhome/production5` directory, the correct permissions are shown, where `production5` is the user owner and the group owner.

5. On `serverb`, adjust the appropriate SELinux Boolean so that the `production5` user may use the NFS-mounted home directory after authenticating with an SSH key. If required, use `redhat` as the password of the `production5` user.
 - 5.1. Open a new terminal window and verify from `servera` that the `production5` user cannot log in to `serverb` with SSH key-based authentication. An SELinux Boolean is

preventing the user from logging in. From `workstation`, open a new terminal and log in to `servera` as the `student` user.

```
[student@workstation ~]$ ssh student@servera
...output omitted...
[student@servera ~]$
```

- 5.2. Switch to the `production5` user. When prompted, use `redhat` as the password of the `production5` user.

```
[student@servera ~]$ su - production5
Password: redhat
[production5@servera ~]$
```

- 5.3. Generate an SSH key pair.

```
[production5@servera ~]$ ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/home/production5/.ssh/id_rsa): Enter
Created directory '/home/production5/.ssh'.
Enter passphrase (empty for no passphrase): Enter
Enter same passphrase again: Enter
Your identification has been saved in /home/production5/.ssh/id_rsa.
Your public key has been saved in /home/production5/.ssh/id_rsa.pub.
The key fingerprint is:
SHA256:AbUcIBXneyiGIhr4wS1xzs3WqDvbTP+eZuSRn9HQ/cw
production5@servera.lab.example.com
The key's randomart image is:
+---[RSA 3072]-----+
| ..=++ |
| . = 0 |
| . . = . . |
|.. * + 0 + . . .|
|+ = = B S .. 0 0.|
|. + + + . . + . . E|
|. . . . 0 0 0 |
| . = . +.0 |
| 000 . = + |
+----[SHA256]-----+
```

- 5.4. Transfer the public key of the SSH key pair to the `production5` user on the `serverb` machine. When prompted, use `redhat` as the password of the `production5` user.

```
[production5@servera ~]$ ssh-copy-id production5@serverb
/usr/bin/ssh-copy-id: INFO: Source of key(s) to be installed: "/home/
production5/.ssh/id_rsa.pub"
The authenticity of host 'serverb (172.25.250.11)' can't be established.
ECDSA key fingerprint is SHA256:ciCkaRWF4g6eR9nSdPxQ7KL8czpViXal6BousK544TY.
Are you sure you want to continue connecting (yes/no)? yes
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter
out any that are already installed
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompted
now it is to install the new keys
```



```
production5@serverb's password: redhat
```

```
Number of key(s) added: 1
```

Now try logging in to the machine, with: "ssh 'production5@serverb'" and check to make sure that only the key(s) you wanted were added.

- 5.5. Use SSH public key-based authentication instead of password-based authentication to log in to `serverb` as the `production5` user. This command should fail.

```
[production5@servera ~]$ ssh -o pubkeyauthentication=yes \
-o passwordauthentication=no production5@serverb
production5@serverb: Permission denied (publickey,gssapi-keyex,gssapi-with-
mic,password).
```

- 5.6. On the terminal that is connected to `serverb` as the `root` user, set the `use_nfs_home_dirs` SELinux Boolean to `true`.

```
[root@serverb ~]# setsebool -P use_nfs_home_dirs true
```

- 5.7. Return to the terminal that is connected to `servera` as the `production5` user, and use SSH public key-based authentication instead of password-based authentication to log in to `serverb` as the `production5` user. This command should succeed.

```
[production5@servera ~]$ ssh -o pubkeyauthentication=yes \
-o passwordauthentication=no production5@serverb
...output omitted...
[production5@serverb ~]$
```

- 5.8. Exit and close the terminal that is connected to `serverb` as the `production5` user. Keep open the terminal that is connected to `serverb` as the `root` user.
6. On `serverb`, adjust the firewall settings to block all connection requests that originate from the `servera` machine. Use the `servera` IPv4 address (`172.25.250.10`) to configure the firewall rule.
 - 6.1. Add the IPv4 address of `servera` to the `block` zone.

```
[root@serverb ~]# firewall-cmd --add-source=172.25.250.10/32 \
--zone=block --permanent
success
```

- 6.2. Reload the changes in the firewall settings.

```
[root@serverb ~]# firewall-cmd --reload
success
```

7. On `serverb`, investigate and fix the issue with the failing Apache web service, which listens on port `30080/TCP` for connections. Adjust the firewall settings appropriately so that the port `30080/TCP` is open for incoming connections.
 - 7.1. Restart the `httpd` service. This command fails to restart the service.

```
[root@serverb ~]# systemctl restart httpd.service
Job for httpd.service failed because the control process exited with error code.
See "systemctl status httpd.service" and "journalctl -xeu httpd.service" for
details.
```

- 7.2. Investigate why the `httpd` service is failing. A permission error indicates that the `httpd` daemon failed to bind to port `30080/TCP` on startup. SELinux policies can prevent an application from binding to a non-standard port. Press `q` to quit the command.

```
[root@serverb ~]# systemctl status httpd.service
× httpd.service - The Apache HTTP Server
   Loaded: loaded (/usr/lib/systemd/system/httpd.service; disabled; vendor
   preset: disabled)
   Active: failed (Result: exit-code) since Mon 2022-05-02 13:20:46 EDT; 29s ago
     Docs: man:httpd.service(8)
   Process: 2322 ExecStart=/usr/sbin/httpd $OPTIONS -DFOREGROUND (code=exited,
   status=1/FAILURE)
   Main PID: 2322 (code=exited, status=1/FAILURE)
     Status: "Reading configuration..."
      CPU: 30ms

May 02 13:20:46 serverb.lab.example.com systemd[1]: Starting The Apache HTTP
   Server...
May 02 13:20:46 serverb.lab.example.com httpd[2322]: (13)Permission denied:
   AH00072: make_sock: could not bind to address [::]:30080
May 02 13:20:46 serverb.lab.example.com httpd[2322]: (13)Permission denied:
   AH00072: make_sock: could not bind to address 0.0.0.0:30080
May 02 13:20:46 serverb.lab.example.com httpd[2322]: no listening sockets
   available, shutting down
...output omitted...
```

- 7.3. Determine whether an SELinux policy is preventing the `httpd` service from binding to the `30080/TCP` port. The log messages reveal that the `30080/TCP` port does not have the appropriate `http_port_t` SELinux context, and so SELinux prevents the `httpd` service from binding to the port. The log message also produces the syntax of the `semanage port` command, so that you can fix the issue.

```
[root@serverb ~]# sealert -a /var/log/audit/audit.log
...output omitted...
SELinux is preventing /usr/sbin/httpd from name_bind access on the tcp_socket port
30080.

***** Plugin bind_ports (92.2 confidence) suggests *****

If you want to allow /usr/sbin/httpd to bind to network port 30080
Then you need to modify the port type.
Do
# semanage port -a -t PORT_TYPE -p tcp 30080
   where PORT_TYPE is one of the following: http_cache_port_t, http_port_t,
   jboss_management_port_t, jboss_messaging_port_t, ntop_port_t, puppet_port_t.
...output omitted...
```

- 7.4. Set the appropriate SELinux context on the 30080/TCP port for the `httpd` service to bind to it.

```
[root@serverb ~]# semanage port -a -t http_port_t -p tcp 30080
```

- 7.5. Restart the `httpd` service. This command should successfully restart the service.

```
[root@serverb ~]# systemctl restart httpd
```

- 7.6. Add the 30080/TCP port to the default `public` zone.

```
[root@serverb ~]# firewall-cmd --add-port=30080/tcp --permanent
success
[root@serverb ~]# firewall-cmd --reload
success
```

- 7.7. Return to the `workstation` machine as the `student` user.

```
[root@serverb ~]# exit
logout
[student@serverb ~]$ exit
logout
Connection to serverb closed.
```

Evaluation

As the `student` user on the `workstation` machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade rhcsa-compreview3
```

Finish

On the `workstation` machine, change to the `student` user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish rhcsa-compreview3
```

► Lab

Run Containers



Note

If you plan to take the RHCSA exam, then use the following approach to maximize the benefit of this Comprehensive Review: attempt each lab without viewing the solution buttons or referring to the course content. Use the grading scripts to gauge your progress as you complete each lab.

Create rootless detached containers.

Outcomes

- Create rootless detached containers.
- Configure port mapping and persistent storage.
- Configure a container as a `systemd` service and use `systemctl` commands to manage it.

Before You Begin

If you did not reset your `workstation` and `server` machines at the end of the last chapter, then save any work that you want to keep from earlier exercises on those machines, and reset them now.

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start rhcsa-compreview4
```

Specifications

- On `serverb`, configure the `podmgr` user with `redhat` as the password, and set up the appropriate tools for the `podmgr` user to manage the containers for this comprehensive review. Configure `registry.lab.example.com` as a remote registry. Use `admin` as the user and `redhat321` as the password to authenticate to the registry. You can use the `/tmp/review4/registries.conf` file to configure the registry.
- The `/tmp/review4/container-dev` directory contains two directories with development files for the containers in this comprehensive review. Copy the two directories under the `/tmp/review4/container-dev` directory to the `podmgr` home directory. Configure the `/home/podmgr/storage/database` subdirectory so that you can use it as persistent storage for a container.
- Create the `db-app01` detached container based on the `registry.lab.example.com/rhel9/mariadb-105` container image. Use the `/home/podmgr/storage/database`

directory as persistent storage for the `/var/lib/mysql/data` directory of the `db-app01` container. Map the 13306 port on the local machine to the 3306 port in the container. Use the values of the following table to set the environment variables to create the containerized database:

Variable	Value
MYSQL_USER	developer
MYSQL_PASSWORD	redhat
MYSQL_DATABASE	inventory
MYSQL_ROOT_PASSWORD	redhat

- Create a `systemd` service file to manage the `db-app01` container. Configure the `systemd` service so that when you start the service, the `systemd` daemon keeps the original container. Start and enable the container as a `systemd` service. Configure the `db-app01` container to start at system boot.
- Copy the `/home/podmgr/db-dev/inventory.sql` script into the `/tmp` directory of the `db-app01` container, and execute the script inside the container. If you executed the script locally, then you would use the `mysql -u root inventory < /tmp/inventory.sql` command.
- Use the container file in the `/home/podmgr/http-dev` directory to create the `http-app01` detached container. The container image name must be `http-client`. Map the 8080 port on the local machine to the 8080 port in the container.
- Use the `curl` command to query the content of the `http-app01` container. Verify that the output of the command shows the container name of the client and that the status of the database is up.

Evaluation

As the student user on the workstation machine, use the `lab` command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade rhcsa-compreview4
```

Finish

On the workstation machine, change to the student user home directory and use the `lab` command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish rhcsa-compreview4
```

► Solution

Run Containers



Note

If you plan to take the RHCSA exam, then use the following approach to maximize the benefit of this Comprehensive Review: attempt each lab without viewing the solution buttons or referring to the course content. Use the grading scripts to gauge your progress as you complete each lab.

Create rootless detached containers.

Outcomes

- Create rootless detached containers.
- Configure port mapping and persistent storage.
- Configure a container as a `systemd` service and use `systemctl` commands to manage it.

Before You Begin

If you did not reset your `workstation` and `server` machines at the end of the last chapter, then save any work that you want to keep from earlier exercises on those machines, and reset them now.

As the `student` user on the `workstation` machine, use the `lab` command to prepare your system for this exercise.

This command prepares your environment and ensures that all required resources are available.

```
[student@workstation ~]$ lab start rhcsa-compreview4
```

1. On `serverb`, configure the `podmgr` user with `redhat` as the password and set up the appropriate tools for the `podmgr` user to manage the containers for this comprehensive review. Configure the `registry.lab.example.com` as the remote registry. Use `admin` as the user and `redhat321` as the password to authenticate. You can use the `/tmp/review4/registries.conf` file to configure the registry.

- 1.1. Log in to `serverb` as the `student` user.

```
[student@workstation ~]$ ssh student@serverb
...output omitted...
[student@serverb ~]$
```

- 1.2. Install the `container-tools` meta-package.

```
[student@serverb ~]$ sudo dnf install container-tools
[sudo] password for student: student
...output omitted...
Is this ok [y/N]: y
...output omitted...
Complete!
```

- 1.3. Create the `podmgr` user and set `redhat` as the password for the user.

```
[student@serverb ~]$ sudo useradd podmgr
[student@serverb ~]$ sudo passwd podmgr
Changing password for user podmgr.
New password: redhat
BAD PASSWORD: The password is shorter than 8 characters
Retype new password: redhat
passwd: all authentication tokens updated successfully.
```

- 1.4. Exit the `student` user session. Log in to the `serverb` machine as the `podmgr` user. If prompted, use `redhat` as the password.

```
[student@serverb ~]$ exit
logout
Connection to serverb closed.
[student@workstation ~]$ ssh podmgr@serverb
...output omitted...
[podmgr@serverb ~]$
```

- 1.5. Create the `~/.config/containers` directory.

```
[podmgr@serverb ~]$ mkdir -p ~/.config/containers
```

- 1.6. Copy the `/tmp/review4/registries.conf` file to the container configuration directory in the home directory.

```
[podmgr@serverb ~]$ cp /tmp/review4/registries.conf ~/.config/containers/
```

- 1.7. Log in to the registry to verify the configuration.

```
[podmgr@serverb ~]$ podman login registry.lab.example.com
Username: admin
Password: redhat321
Login Succeeded!
```

2. The `/tmp/review4/container-dev` directory contains two directories with development files for the containers in this comprehensive review. Copy the two directories in the `/tmp/review4/container-dev` directory to the `podmgr` home directory. Configure the `/home/podmgr/storage/database` subdirectory so that you can use it as persistent storage for a container.
 - 2.1. Copy the content of the `/tmp/review4/container-dev` directory to the `podmgr` home directory.

```
[podmgr@serverb ~]$ cp -r /tmp/review4/container-dev/*-dev .
[podmgr@serverb ~]$ ls -l
total 0
drwxr-xr-x. 2 podmgr podmgr 27 Jun 18 12:00 db-dev
drwxr-xr-x. 3 podmgr podmgr 42 Jun 18 12:00 http-dev
```

- 2.2. Create the `/home/podmgr/storage/database` directory in the podmgr home directory. Set the appropriate permissions on the directory for the container to mount it as persistent storage.

```
[podmgr@serverb ~]$ mkdir -p storage/database
[podmgr@serverb ~]$ chmod 0777 storage/database
[podmgr@serverb ~]$ ls -l storage/
total 0
drwxrwxrwx. 2 podmgr podmgr 6 Jun 18 12:05 database
```

- 3. Create the `db-app01` detached container based on the `registry.lab.example.com/rhel9/mariadb-105` container image. Use the `/home/podmgr/storage/database` directory as persistent storage for the `/var/lib/mysql/data` directory of the `db-app01` container. Map the 13306 port on the local machine to the 3306 port in the container. Use the values of the following table to set the environment variables to create the containerized database:

Variable	Value
MYSQL_USER	developer
MYSQL_PASSWORD	redhat
MYSQL_DATABASE	inventory
MYSQL_ROOT_PASSWORD	redhat

- 3.1. Inspect the `registry.lab.example.com/rhel9/mariadb` container image.

```
[podmgr@serverb ~]$ skopeo inspect \
docker://registry.lab.example.com/rhel9/mariadb-105
{
  "Name": "registry.lab.example.com/rhel9/mariadb-105",
  "Digest": "sha256:2bf8...fab8",
  "RepoTags": [
    "latest"
  ],
  ...output omitted...
```

- 3.2. Use the `registry.lab.example.com/rhel9/mariadb-105` container image to create the detached `db-app01` container. Use the `/home/podmgr/storage/database` directory as persistent storage for the container. Map the 13306 port to the 3306 container port. Use the data in the table to set the environment variables for the container.


```
[podmgr@serverb ~]$ podman run -d --name db-app01 \
-e MYSQL_USER=developer \
-e MYSQL_PASSWORD=redhat \
-e MYSQL_DATABASE=inventory \
-e MYSQL_ROOT_PASSWORD=redhat \
-p 13306:3306 \
-v /home/podmgr/storage/database:/var/lib/mysql/data:Z \
registry.lab.example.com/rhel9/mariadb-105
...output omitted...
0959ba5df57dec2a478504df633a8cdb54aaad5c98a160de34d1d1b65f8435a6
[podmgr@serverb ~]$ podman ps -a
```

CONTAINER ID	IMAGE	COMMAND
0959ba5df57d	registry.lab.example.com/rhel9/mariadb-105:latest	run-mysqld 16

```
seconds ago Up 15 seconds 0.0.0.0:13306->3306/tcp db-app01
```

4. Create a systemd service file to manage the db-app01 container. Configure the systemd service so that when you start the service, the systemd daemon keeps the original container. Start and enable the container as a systemd service. Configure the db-app01 container to start at system boot.

- 4.1. Create the ~/.config/systemd/user/ directory for the container unit file.

```
[podmgr@serverb ~]$ mkdir -p ~/.config/systemd/user/
```

- 4.2. Create the systemd unit file for the db-app01 container, and move the unit file to the ~/.config/systemd/user/ directory.

```
[podmgr@serverb ~]$ podman generate systemd --name db-app01 --files
/home/podmgr/container-db-app01.service
[podmgr@serverb ~]$ mv container-db-app01.service ~/.config/systemd/user/
```

- 4.3. Stop the db-app01 container.

```
[podmgr@serverb ~]$ podman stop db-app01
db-app01
[podmgr@serverb ~]$ podman ps -a
```

CONTAINER ID	IMAGE	COMMAND
0959ba5df57d	registry.lab.example.com/rhel9/mariadb-105:latest	run-mysqld 2

```
minutes ago Exited (0) 2 seconds ago 0.0.0.0:13306->3306/tcp db-app01
```

- 4.4. Reload the user systemd service to use the new service unit.

```
[podmgr@serverb ~]$ systemctl --user daemon-reload
```

- 4.5. Start and enable the systemd unit for the db-app01 container.

```
[podmgr@serverb ~]$ systemctl --user enable --now container-db-app01
Created symlink /home/podmgr/.config/systemd/user/default.target.wants/container-
db-app01.service → /home/podmgr/.config/systemd/user/container-db-app01.service.
[podmgr@serverb ~]$ systemctl --user status container-db-app01
● container-db-app01.service - Podman container-db-app01.service
   Loaded: loaded (/home/podmgr/.config/systemd/user/container-db-app01.service;
   enabled; preset: disabled)
   Active: active (running) since Fri 2024-06-18 12:18:45 EDT; 13s ago
...output omitted...
[podmgr@serverb ~]$ podman ps -a
```

CONTAINER ID	IMAGE	COMMAND
0959ba5df57d	registry.lab.example.com/rhel9/mariadb-105:latest	run-mysqld 2
minutes ago	Up 2 minutes	0.0.0.0:13306->3306/tcp db-app01

- 4.6. Use the `loginctl` command to configure the `db-app01` container to start at system boot.

```
[podmgr@serverb ~]$ loginctl enable-linger
```

5. Copy the `/home/podmgr/db-dev/inventory.sql` script into the `/tmp` directory of the `db-app01` container, and execute the script inside the container. If you executed the script locally, then you would use the `mysql -u root inventory < /tmp/inventory.sql` command.
 - 5.1. Copy the `/home/podmgr/db-dev/inventory.sql` script into the `/tmp` directory of the `db-app01` container.

```
[podmgr@serverb ~]$ podman cp /home/podmgr/db-dev/inventory.sql \
db-app01:/tmp/inventory.sql
```

- 5.2. Execute the `inventory.sql` script in the `db-app01` container.

```
[podmgr@serverb ~]$ podman exec -it db-app01 sh -c 'mysql -u root \
inventory < /tmp/inventory.sql'
```

6. Use the container file in the `/home/podmgr/http-dev` directory to create the `http-app01` detached container. The container image name must be `http-client`. Map the 8080 port on the local machine to the 8080 port in the container.
 - 6.1. Create the `http-client:9.0` image with the container file in the `/home/podmgr/http-dev` directory.

```
[podmgr@serverb ~]$ podman build -t http-client:9.0 http-dev/
STEP 1/7: FROM registry.lab.example.com/rhel9/php-82:1-15
...output omitted...
```

- 6.2. Create the `http-app01` detached container. Map the 8080 port from the local machine to the 8080 port in the container.

```
[podmgr@serverb ~]$ podman run -d --name http-app01 \
-p 8080:8080 localhost/http-client:9.0
[podmgr@serverb ~]$ podman ps -a
```

CONTAINER ID	IMAGE	STATUS	PORTS	NAMES	COMMAND	CREATED
ba398d080e00	registry.lab.example.com/rhel9/mariadb-105:1-86	Up 20 minutes ago	0.0.0.0:13306->3306/tcp	db-app01	run-mysqld	20
ee424df19621	localhost/http-client:9.0	Up 4 seconds ago	0.0.0.0:8080->8080/tcp	http-app01	/bin/sh -c	4

- 7. Query the content of the http-app01 container. Verify that it shows the container name of the client and that the status of the database is up.
 - 7.1. Verify that the http-app01 container responds to http requests.

```
[podmgr@serverb ~]$ curl 127.0.0.1:8080
This is the server http-app01 and the database is up
```

- 8. Return to the workstation machine as the student user.

```
[podmgr@serverb ~]$ exit
logout
Connection to serverb closed.
[student@workstation ~]$
```

Evaluation

As the student user on the workstation machine, use the lab command to grade your work. Correct any reported failures and rerun the command until successful.

```
[student@workstation ~]$ lab grade rhcsa-compreview4
```

Finish

On the workstation machine, change to the student user home directory and use the lab command to complete this exercise. This step is important to ensure that resources from previous exercises do not impact upcoming exercises.

```
[student@workstation ~]$ lab finish rhcsa-compreview4
```