

Introduction

1. Red Hat System Administration I

Red Hat System Administration I introduction

Red Hat System Administration I (RH124) is designed for IT professionals without previous Linux system administration experience. The course is intended to provide students with **Linux administration** "survival skills" by focusing on **core administration tasks**. *Red Hat System Administration I* also provides a foundation for students planning to become **full-time Linux system administrators** by introducing key **command-line** concepts and **enterprise-level tools**. These concepts are further developed in the follow-on course, *Red Hat System Administration II (RH134)*.

生存技巧

Course objectives

- 获得足够的技巧去完成系统的核心管理任务
- Gain sufficient skill to perform core system administration tasks on Red Hat Enterprise Linux.
- Build foundational skills needed by an **RHCSA**-certified Red Hat Enterprise Linux system administrator.

Audience

- IT professionals across a broad range of disciplines who need to perform essential Linux administration tasks, including **installation**, **establishing network connectivity**, **managing physical storage** and **basic security administration**.

Prerequisites

- 无需任何基础
- There are **no formal prerequisites for this course**; however, **previous system administration experience on other operating systems will be very beneficial**.

2. Orientation to the Classroom Environment

In this course, students will do most hands-on practice exercises and lab work with two computer systems, which will be referred to as **desktop** and **server**. These machines have the host names `desktopX.example.com` and `serverX.example.com`, where the *x* in the computers' host names will be a number that will vary from student to student. Both machines have a standard user account, *student*, with the password *student*. The *root* password on both systems is *redhat*. Each student is on the IPv4 network `172.25.x.0/24`, where the *x* matches the number of their `desktopX` and `serverX` systems. The instructor runs a central utility server, `classroom.example.com`, which acts as a router for the classroom networks and provides DNS, DHCP, HTTP, and other content services.

□

Classroom Machines

Machine name	IP addresses	Role
<code>desktopX.example.com</code>	<code>172.25.x.10</code>	Student "client" computer

Machine name	IP addresses	Role
serverX.example.com	172.25.x.11	Student "server" computer
classroom.example.com	172.25.254.254	Classroom utility server

3. Internationalization

Language support

Red Hat Enterprise Linux 7 officially supports **22 languages**: English, Assamese, Bengali, Chinese (Simplified), Chinese (Traditional), French, German, Gujarati, Hindi, Italian, Japanese, Kannada, Korean, Malayalam, Marathi, Odia, Portuguese (Brazilian), Punjabi, Russian, Spanish, Tamil, and Telugu.

Per-user language selection

Users may prefer to use a different language for their desktop environment than the system-wide default. They may also want to set their account to use a different keyboard layout or input method.

Language settings

In the GNOME desktop environment, the user may be prompted to set their preferred language and input method on first login. If not, then the easiest way for an individual user to **adjust** their preferred language and input method settings is to use the **Region & Language** application. Run the command **gnome-control-center region**, or from the top bar, select (User) → **Settings**. In the window that opens, select **Region & Language**. The user can click the **Language** box and select their preferred language from the list that appears. This will also update the **Formats** setting to the default for that language. The next time the user logs in, these changes will take full effect.

调节

These settings affect the GNOME desktop environment and any applications,

通过gnome-control-center的设置仅会影响到GNOME桌面环境和相关的应用，但不会应用到SSH、TTY登录

including **gnome-terminal**, started inside it. However, they do not apply to that account if accessed through an **ssh** login from a remote system or a local text console (such as **tty2**).

Note

A user can make their shell environment use the same **LANG** setting as their graphical environment, **even** when they log in through a text console or over **ssh**. One way to do this is to place code similar to the following in the user's **~/.bashrc** file. This example code will set the language used on a text login to match the one currently set for the user's GNOME desktop environment:

一致的

```
i=$(grep 'Language=' /var/lib/AccountService/users/${USER} \
| sed 's/Language=//')
if [ "$i" != "" ]; then
    export LANG=$i
fi
```

已接受

Japanese, Korean, Chinese, or other languages with a non-Latin character set may not display properly on local text consoles.

Individual commands can be made to use another language by setting the **LANG** variable on the command line:

```
[user@host ~]$ LANG=fr_FR.utf8 date  
jeu. avril 24 17:55:01 CDT 2014
```

Subsequent commands will revert to using the system's default language for output.

后来的命令又恢复到系统默认的语言设置

The **locale** command can be used to check the current value of **LANG** and other related environment variables.

Input method settings

IBus 全称 Intelligent Input Bus 是下一代输入法框架（或者说“平台”）。项目现托管于 Google Code - <https://code.google.com/p/ibus/> 此项目包含了世界多数语言的文字输入需求——由世界多个国家开发者维护。

GNOME 3 in Red Hat Enterprise Linux 7 automatically uses the **IBus** input method selection system, which makes it easy to change keyboard layouts and input methods quickly.

<http://wiki.ubuntu.org.cn/ibus>

The **Region & Language** application can also be used to enable alternative input methods. In the **Region & Language** application's window, the **Input Sources** box shows what input methods are currently available. By default, **English (US)** may be the only available method. Highlight **English (US)** and click the **keyboard** icon to see the current keyboard layout.

To add another input method, click the **+** button at the bottom left of the **Input Sources** window. An **Add an Input Source** window will open. Select your language, and then your preferred input method or keyboard layout.

Once more than one input method is configured, the user can switch between them quickly by typing **Super+Space** (sometimes called **Windows+Space**). A *status indicator* will also appear in the GNOME top bar, which has two functions: It indicates which input method is active, and acts as a menu that can be used to switch between input methods or select advanced features of more complex input methods.

Some of the methods are marked with gears, which indicate that those methods have advanced configuration options and capabilities. For example, the Japanese **Japanese (Kana Kanji)** input method allows the user to pre-edit text in Latin and use **Down Arrow** and **Up Arrow** keys to select the correct characters to use.

US English speakers may find also this useful. For example, under **English (United States)** is the keyboard layout **English (international AltGr dead keys)**, which treats **AltGr** (or the right **Alt**) on a PC 104/105-key keyboard as a "secondary-shift" modifier key and dead key activation key for typing additional characters. There are also Dvorak and other alternative layouts available.

Note

Any Unicode character can be entered in the GNOME desktop environment if the user knows the character's Unicode code point, by typing **Ctrl+Shift+U**, followed by the code point.

After **Ctrl+Shift+U** has been typed, an underlined **u** will be displayed to indicate that the system is waiting for Unicode code point entry.

For example, the lowercase Greek letter lambda has the code point U+03BB, and can be entered by typing **Ctrl+Shift+U**, then **03bb**, then **Enter**.

System-wide default language settings

整个系统的默认语言设置

The system's default language is set to US English, using the UTF-8 encoding of Unicode as its character set (**en_US.utf8**), but this can be changed during or after installation.

输入模式设置

From the command line, *root* can change the system-wide locale settings with the `localectl` command. If `localectl` is run with no arguments, it will display the current system-wide locale settings.

To set the system-wide language, run the command `localectl set-locale LANG=locale`, where *locale* is the appropriate `$LANG` from the "Language Codes Reference" table in this chapter. The change will take effect for users on their next login, and is stored in `/etc/locale.conf`.

```
[root@host ~]# localectl set-locale LANG=fr_FR.utf8
```

In GNOME, an administrative user can change this setting from **Region & Language** and clicking the **Login** screen button at the upper-right corner of the window. Changing the **Language** of the login screen will also adjust the system-wide default language setting stored in the `/etc/locale.conf` configuration file.

Important

Local text consoles such as `tty2` are more limited in the fonts that they can display than `gnome-terminal` and `ssh` sessions. For example, Japanese, Korean, and Chinese characters may not display as expected on a local text console. For this reason, it may make sense to use English or another language with a Latin character set for the system's text console.

Likewise, local text consoles are more limited in the input methods they support, and this is managed separately from the graphical desktop environment. The available global input settings can be configured through `localectl` for both local text virtual consoles and the X11 graphical environment. See the `localectl(1)`, `kbd(4)`, and `vconsole.conf(5)` man pages for more information.

建议在文本模式下使用英文语言环境

Language packs

When using non-English languages, you may want to install additional "language packs" to provide additional translations, dictionaries, and so forth. To view the list of available langpacks, run `yum langavailable`. To view the list of langpacks currently installed on the system, run `yum langlist`. To add an additional langpack to the system, run `yum langinstall code`, where *code* is the code in square brackets after the language name in the output of `yum langavailable`.

References

`locale(7)`, `localectl(1)`, `kbd(4)`, `locale.conf(5)`, `vconsole.conf(5)`, `unicode(7)`, `utf-8(7)`, and `yum-langpacks(8)` man pages

Conversions between the names of the graphical desktop environment's X11 layouts and their names in `localectl` can be found in the file `/usr/share/X11/xkb/rules/base.lst`.

3.1. Language Codes Reference

□

Language Codes

Language	\$LANG value
English (US)	en_US.utf8
Assamese	as_IN.utf8

Language	\$LANG value
Bengali	bn_IN. utf8
Chinese (Simplified)	zh_CN. utf8
Chinese (Traditional)	zh_TW. utf8
French	fr_FR. utf8
German	de_DE. utf8
Gujarati	gu_IN. utf8
Hindi	hi_IN. utf8
Italian	it_IT. utf8
Japanese	ja_JP. utf8
Kannada	kn_IN. utf8
Korean	ko_KR. utf8
Malayalam	ml_IN. utf8
Marathi	mr_IN. utf8
Odia	or_IN. utf8
Portuguese (Brazilian)	pt_BR. utf8
Punjabi	pa_IN. utf8
Russian	ru_RU. utf8
Spanish	es_ES. utf8
Tamil	ta_IN. utf8
Telugu	te_IN. utf8

▮ Chapter 1. Accessing the Command Line

1.1. Accessing the Command Line Using the Local Console

1.2. Practice: Local Console Access Terms

1.3. Accessing the Command Line Using the Desktop

1.4. Practice: The GNOME 3 Desktop Environment

1.5. Executing Commands Using the Bash Shell

1.6. Practice: Bash Commands and Keyboard Shortcuts

1.7. Lab: Accessing the Command Line

Introduction

Overview	
Goal	To log into a Linux system and run simple commands using the shell.
Objectives	<ul style="list-style-type: none">Use Bash shell syntax to enter commands at a Linux console.Launch applications in a GNOME desktop environment.Use Bash features to run commands from a shell prompt using fewer keystrokes. 快捷键的使用
Sections	<ul style="list-style-type: none">Accessing the Command Line Using the Local Console (and Practice)Accessing the Command Line Using the Desktop (and Practice)Executing Commands Using the Bash Shell (and Practice)
Lab	<ul style="list-style-type: none">Accessing the Command Line

1.1. Accessing the Command Line Using the Local Console

Objectives

After completing this section, students should be able to log into a Linux system on a local text console and run simple commands using the shell.

The **bash** shell

Accessing the command line with virtual consoles

A *command line* is a text-based interface which can be used to input instructions to a computer system.

The Linux command line is provided by a program called the *shell*. Over the long history of UNIX-like systems, many shells have been developed. The default shell for users in Red Hat Enterprise Linux is the **GNU Bourne-Again Shell (bash)**. Bash is an improved version of one of the most successful shells used on UNIX-like systems, the **Bourne Shell (sh)**.

When a shell is used interactively, it displays a string when it is waiting for a command from the user.

This is called the *shell prompt*. When a regular user starts a shell, the default prompt ends with

a **\$** character.

```
[student@desktopX ~]$
```

The **\$** is replaced by a **#** if the shell is running as the superuser, **root**. This makes it more obvious that it is a superuser shell, which helps to avoid accidents and mistakes in the privileged account.

```
[root@desktopX ~]#
```

Using **bash** to execute commands can be **powerful**. The **bash** shell provides a **scripting language that can support automation of tasks**. The shell has additional capabilities that can simplify or make possible operations that are hard to accomplish **efficiently** with graphical tools. 强大、高效、自动

Note

The **bash** shell is similar in concept to the command line interpreter found in recent versions of Microsoft Windows **cmd.exe**, although **bash** has a more sophisticated scripting language. It is also similar to Windows PowerShell in Windows 7 and Windows Server 2008 R2. Mac OS X administrators who use the Macintosh's Terminal utility may be pleased to note that **bash** is the default shell in Mac OS X.

Virtual consoles 虚拟控制台

Users access the **bash** shell through a *terminal*. A terminal provides a keyboard for user input and a display for output. On text-based installations, this can be the Linux machine's *physical console*, the hardware keyboard and display. Terminal access can also be configured through serial ports.

Another way to access a shell is from a *virtual console*. A Linux machine's physical console supports multiple virtual consoles which act like separate terminals. Each virtual console supports an independent login session. 每个虚拟的控制台就像一个分离的终端一样，支持一个独立的会话

If the graphical environment is available, it will run on the *first* virtual console in Red Hat Enterprise Linux 7. Five additional text login prompts are available on consoles two through six (or one through five if the graphical environment is turned off). With a graphical environment running, access a text login prompt on a virtual console by holding **Ctrl+Alt** and pressing a function key (**F2** through **F6**). Press **Ctrl+Alt+F1** to return to the first virtual console and the graphical desktop.

Important

In the pre-configured virtual images delivered by Red Hat, login prompts have been disabled in the virtual consoles.

Note

In Red Hat Enterprise Linux 5 and earlier, the first six virtual consoles always provided text login prompts. When the graphical environment was launched, it ran on virtual console seven (accessed through **Ctrl+Alt+F7**).

Shell basics

Commands entered at the shell prompt have three basic parts:

- *Command* to run
- **Options** to adjust the behavior of the command
- **Arguments** which are typically targets of the command

选项改变命令的行为
参数指出命令操作的目标

The *command* is the name of the program to run. It may be followed by one or more *options*, which adjust the behavior of the command or what it will do. Options normally start with one or two dashes (**-a** or **--all**, for example) to distinguish them from arguments. Commands may also be followed by one or more *arguments*, which often indicate a target that the command should operate on. 一个或两个破折号

For example, the command line **usermod -L morgan** has a command (**usermod**), an option (**-L**), and an argument (**morgan**). The effect of this command is to lock the password on user morgan's account.

To use a command effectively, a user needs to know what options and arguments it takes and in what order it expects them (the *syntax* of the command). Most commands have a **--help** option. This causes the command to print a description of what it does, a "usage statement" that describes the command's syntax, and a list of the options it accepts and what they do.

Usage statements may seem complicated and difficult to read. They become much simpler to understand once a user becomes familiar with a few basic conventions: 一些基本的约定

- Square brackets, **[]**, surround optional items.
- Anything followed by **...** represents an arbitrary-length list of items of that type.

表示一个任意长度的列表

- Multiple items separated by pipes, **|**, means only *one* of them can be specified.
- Text in angle brackets, **<>**, represents variable data. For example, **<filename>** means “insert the filename you wish to use here”. Sometimes these variables are simply written in capital letters (e.g., FILENAME).

Consider the first usage statement for the **date** command:

```
[student@desktopX ~]$ date --help
date [OPTION]... [+FORMAT]
```

显示 This indicates that **date** can take an optional list of options (**[OPTION] . . .**), followed by an optional format string, prefixed with a plus character, **+**, that defines how the current date should be displayed (**[+FORMAT]**). Since both of these are optional, **date** will work even if it is not given options or arguments (it will print the current date and time using its default format). **因为这两个是可选的**

Note

The **man** page for a command has a SYNOPSIS section that provides information about the command's syntax. The **man-pages(7)** man page describes how to interpret all the square brackets, vertical bars, and so forth that users see in SYNOPSIS or a usage message.

When a user is finished using the shell and wants to quit, there are a couple of ways to end the session. The **exit** command terminates the current shell session. Another way to finish a session is by typing **Ctrl+D**.

References

intro(1), **bash(1)**, **console(4)**, **pts(4)**, and **man-pages(7)** man pages
Note: Some details of the console(4) man page, involving init(8) and inittab(5), are outdated.

1.2. Practice: Local Console Access Terms

Quiz **随堂测验**

Match the following items to their counterparts in the table.

Description	Term
The interpreter that executes commands typed as strings.	Shell
The visual cue that indicates an interactive shell is waiting for the user to type a command.	Prompt
The name of a program to run.	Command
The part of the command line that adjusts the behavior of a command.	Option
The part of the command line that specifies the target that the command should operate on.	Argument
The hardware display and keyboard used to interact with a system.	Physical console
One of multiple logical consoles that can each support an independent login session.	Virtual console
An interface that provides a display for output and a keyboard for input to a shell session.	Terminal

1.3. Accessing the Command Line Using the Desktop

Objectives

After completing this section, students should be able to log into the Linux system using the GNOME 3 desktop environment to run commands from a shell prompt in a terminal program.

The GNOME desktop environment

Accessing the command line using the desktop

The *desktop environment* is the graphical user interface on a Linux system. The default desktop environment in Red Hat Enterprise Linux 7 is provided by **GNOME 3**. It provides an integrated desktop for users and a unified development platform on top of a graphical framework provided by the **X Window System**.

The **GNOME Shell** provides the core user interface functions for the GNOME desktop environment. The **gnome-shell** application is highly customizable. By default, RHEL 7 users use the "GNOME Classic" theme for **gnome-shell**, which is similar to the GNOME 2 desktop environment. Another available option is the "modern" GNOME 3 theme used by the upstream GNOME project. Either theme can be selected persistently at login by selecting the gear icon next to the **sign In** button when entering the user's password.

The first time a new user logs in, an initial setup program runs to help them configure basic account settings. The **GNOME Help** application is then started on the **Getting Started with GNOME** screen. This screen includes videos and documentation to help orient new users to the GNOME 3 environment. **GNOME Help** can be quickly started by typing **F1** in **gnome-shell**, by selecting **Applications** → **Documentation** → **Help**, or by running the **yelp** command.




Figure 1.1: An empty GNOME 3 desktop

Parts of the GNOME Shell

The various parts of the GNOME Shell have specific names and purposes. These parts include the following:

- *top bar*: The bar that runs along the top of the screen. The top bar provides the **Applications** and **Places** menus, and controls for volume, networking, calendar access, and selecting between keyboard input methods (if there are more than one configured). Under the menu for the user's name are options to adjust account settings, lock the screen, switch users, log out of the system, or shut it down.
- *Applications menu*: This menu on the top bar provides a way to start applications, categorized into submenus. The **Activities Overview** can also be started from this menu.
- *Places menu*: This menu to the right of the **Applications** menu provides quick access through a graphical file manager to important menus in the user's home directory, to **/**, and to exports and file shares on the network.
- *window list*: The bar that runs along the bottom of the screen. The window list provides an easy way to access, minimize, and restore all windows in the current workspace. On the right corner is an indicator to tell the user which workspace they are on and how many are available.
- *message tray*: The message tray provides a way to review notifications sent by applications or system components to GNOME. If a notification occurs, normally the notification first appears briefly as a single line at the bottom of the screen, and a persistent indicator appears in the lower right corner to inform the user of how many notifications have been recently received. The message tray can be opened to review these notifications by clicking the indicator or typing **Super+M**. The message tray can be closed by typing either **Esc** or **Super+M** again.
- *Activities Overview*: This is a special mode that helps a user organize windows and start applications. The Activities Overview can be started by selecting **Applications** → **Activities Overview**. A faster way is to press the **Super** key (sometimes called the **Windows** key), found near the lower left corner of an IBM PC 104/105-key keyboard. The three main areas of the Activities Overview are the **dash** on the left side of the screen, the **windows overview** in the center of the screen, and the **workspace selector** on the right side of the screen. The Activities Overview can be exited by pressing the **Esc** key.
- *dash*: This is a configurable list of icons of the user's favorite applications, applications which are currently running, and a **grid** button which can be used to select arbitrary applications. Applications can be started by clicking on one of the icons or by using the grid button to find a less commonly used application. The dash is also sometimes called the **dock**.

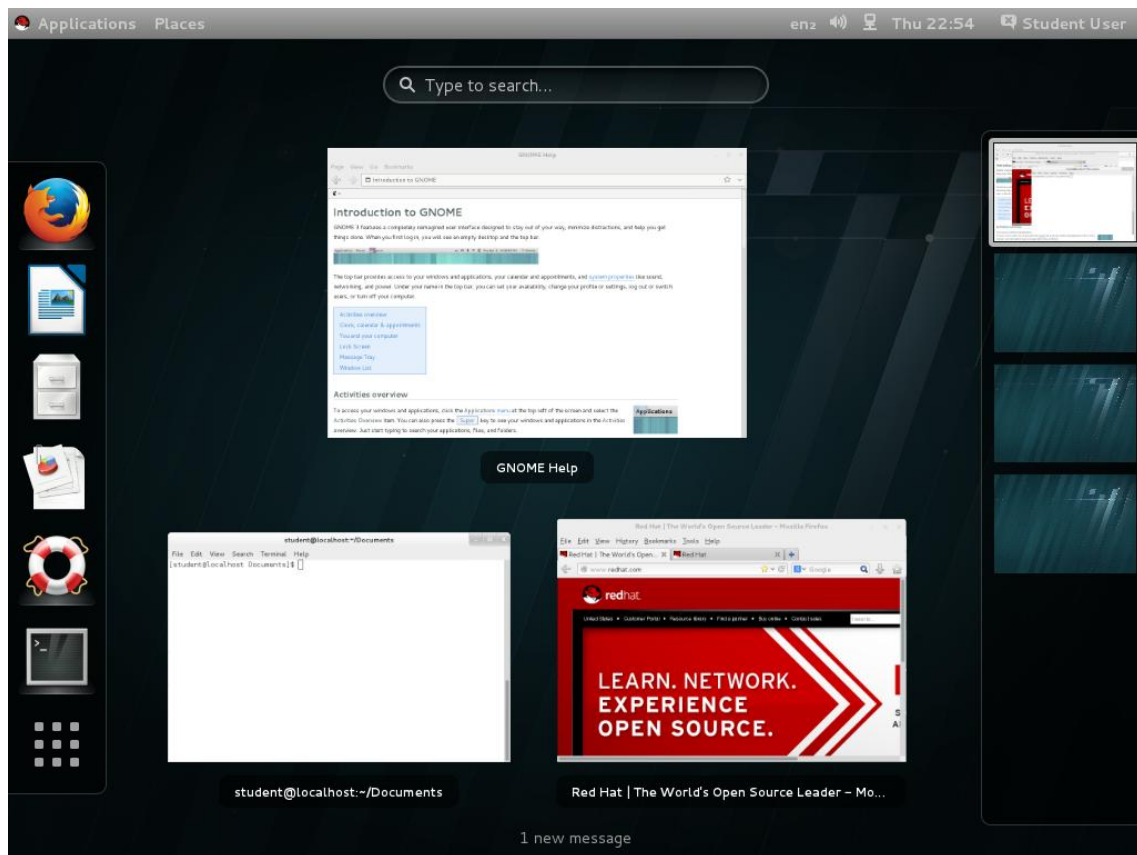


Figure 1.2: The GNOME 3 Activities Overview

Workspaces

Workspaces are separate desktop screens which have different application windows. These can be used to organize the working environment by grouping open application windows by task. For example, windows being used to perform a particular system maintenance activity (such as setting up a new remote server) can be grouped in one workspace, while email and other communication applications can be grouped in another workspace.

There are three methods for switching between workspaces. One method is to click the indicator in the right corner of the window list and select the desired workspace. Another, perhaps the fastest, is to type **Ctrl+Alt+UpArrow** or **Ctrl+Alt+DownArrow** to switch between workspaces sequentially. A third is to switch to the **Activities Overview** and click the desired workspace.

An advantage of using the **Activities Overview** is that windows can be clicked and dragged between the current workspace and one of the others by using the **workspace selector** on the right side of the screen and the **windows overview** in the center of the screen.

Note

Using the **Ctrl+Alt+UpArrow** or **Ctrl+Alt+DownArrow** key combinations to change workspaces does not work in the virtual learning environment. Instead, the workspace applet on the panel or the **Activities Overview** must be used to change workspaces.

Starting a terminal

To get a shell prompt in GNOME, start a graphical terminal application such as **GNOME Terminal**. There are several ways to do this. Here are the three most commonly used methods:

- Select **Applications** → **Utilities** → **Terminal**.
- On an empty desktop, right-click, or press the **Menu** key, and select **Open in Terminal** from the context menu that appears.
- From the **Activities Overview**, select **Terminal** from the **dash** (either from the favorites area or by finding it with either the **grid** button (inside **Utilities** grouping) or the search field at the top of the **windows overview**).

When a terminal window is opened, a shell prompt displays for the user that started the graphical terminal program. The shell prompt and the terminal window's title bar will indicate the current user name, host name, and working directory.

▮ Locking the screen or logging out

Locking the screen, or logging out entirely, can be done from the menu for the user's name on the far right side of the top bar.

To lock the screen, select **(User)** → **Lock** or type **Ctrl+Alt+L**. The screen will lock if the graphical session is idle for a few minutes.

A **lock screen curtain** will appear that shows the system time and the name of the logged-in user. To unlock the screen, press **Enter** or **Space** to raise the lock screen curtain, then enter the user's password on the **lock screen**.

To log out and end the current graphical login session, select **(User)** → **Log Out** from the top bar. A dialog window will appear, giving the option to **Cancel** the log out within 60 seconds, or confirm the **Log Out** action.

▮ Powering off or rebooting the system

To shut down the system, select **(User)** → **Power Off** from the top bar or type **Ctrl+Alt+Del**. In the dialog that appears, the user can choose to **Power Off**, **Restart** the machine, or **Cancel** the operation. If the user does not make a choice in this dialog, the system will automatically shut down after 60 seconds.

References

GNOME Help

- **yelp**

GNOME Help: *Getting Started with GNOME*

- **yelp help:gnome-help/getting-started**

▮ 1.4. Practice: The GNOME 3 Desktop Environment

Guided exercise

In this lab, you will log in through the graphical display manager as a regular user to become familiar with the GNOME Classic desktop environment provided by GNOME 3.

Outcome:

A basic orientation to the GNOME 3 desktop environment.

Before you begin...

Access the graphical login screen of **desktopX.example.com**.

Important

There are two virtual machines available for lab exercises, a desktop machine (generically called **desktopX**) and a server (generically called **serverX**).

Take care to keep straight which virtual machine an exercise wants you to use.

Do each of the following tasks on the **desktopX** machine. Mark each task as it is completed.

1. Log in as **student** using the password **student**.
 - a. At the GNOME login screen, click the **student** user account.
Enter **student** when prompted for the password.
 - b. Click **sign In** once the password has been typed in.
2. Change the password for **student** from **password** to **55TurnK3y**.
 - a. The simplest approach is to open **GNOME Terminal** and use the **passwd** command at the shell prompt.

On the empty desktop, press the **Menu** key or right-click with the mouse to open the context menu.
 - b. Select **Open in Terminal**.
 - c. In the terminal window that appears, type **passwd** at the shell prompt. Follow the instructions provided by the program to change the **student** password from **student** to **55TurnK3y**.
3. Log out.
 - a. Select the **student** → **Log Out** menu item.
 - b. Click the **Log Out** button in the confirmation window that appears.
4. Log back in as **student** with the new password of **55TurnK3y**.
 - a. At the GNOME login screen, click the **student** user account.
Enter **55TurnK3y** when prompted for the password.
 - b. Click **sign In** once the password has been typed in.
5. Lock the screen.
 - a. Select the **student** → **Lock** menu item.
6. Unlock the screen.
 - a. Press **Enter** to lift the lock screen curtain.
 - b. In the **Password** field, enter **55TurnK3y** as the password.
 - c. Click the **Unlock** button. The GNOME desktop should reappear.
7. Determine how to shut down **desktopX** from the graphical interface, but **Cancel** the operation without shutting down the system.
 - a. Select the **student** → **Power Off** menu item.
 - b. Click the **Cancel** button in the confirmation screen that appears.

1.5. Executing Commands Using the Bash Shell

Objectives

After completing this section, students should be able to save time running commands from a shell prompt using Bash shortcuts.

Basic command syntax

Executing Commands Using the Bash Shell

The **GNU Bourne-Again Shell (bash)** is a program that interprets commands typed in by the user. Each string typed into the shell can have up to three parts: the command, options (that begin with a `-` or `--`), and arguments. Each word typed into the shell is separated from each other with spaces. Commands are the names of programs that are installed on the system. Each command has its own options and arguments.

The **Enter** key is pressed when a user is ready to execute a command. Each command is typed on a separate line and the output from each command displays before the shell displays a prompt. If a user wants to type more than one command on a single line, a semicolon, `;`, can be used as a command separator. A semicolon is a member of a class of characters called *metacharacters* that has special meanings for **bash**.

Examples of simple commands

The **date** command is used to display the current date and time. It can also be used by the superuser to set the system clock. An argument that begins with a plus sign (+) specifies a format string for the date command.

```
[student@desktopX ~]$ date
Sat Apr  5 08:13:50 PDT 2014
[student@desktopX ~]$ date +%R
08:13
[student@desktopX ~]$ date +%x
04/05/2014
```

The **passwd** command changes a user's own password. The original password for the account must be specified before a change will be allowed. By default, **passwd** is configured to require a strong password, consisting of lowercase letters, uppercase letters, numbers, and symbols, and is not based on a dictionary word. The superuser can use the **passwd** command to change other users' passwords.

```
[student@desktopX ~]$ passwd
Changing password for user student.
Changing password for student.
(current) UNIX password: old_password
New password: new_password
Retype new password: new_password
```

```
passwd: all authentication tokens updated successfully.
```

Linux does not require file name extensions to classify files by type. The **file** command scans the beginning of a file's contents and displays what type it is. The files to be classified are passed as arguments to the command.

```
[student@desktopX ~]$ file /etc/passwd
/etc/passwd: ASCII text

[student@desktopX ~]$ file /bin/passwd
/bin/passwd: setuid ELF 64-bit LSB shared object, x86-64, version 1
(SYSV), dynamically linked (uses shared libs), for GNU/Linux 2.6.32,
BuildID[sha1]=0x91a7160a019b7f5f754264d920e257522c5bce67, stripped

[student@desktopX ~]$ file /home
/home: directory
```

The **head** and **tail** commands display the beginning and end of a file respectively. By default, these commands display 10 lines, but they both have a **-n** option that allows a different number of lines to be specified. The file to display is passed as an argument to these commands.

```
[student@desktopX ~]$ head /etc/passwd
root:x:0:0:root:/root:/bin/bash
bin:x:1:1:bin:/bin:/sbin/nologin
daemon:x:2:2:daemon:/sbin:/sbin/nologin
adm:x:3:4:adm:/var/adm:/sbin/nologin
lp:x:4:7:lp:/var/spool/lpd:/sbin/nologin
sync:x:5:0:sync:/sbin:/bin/sync
shutdown:x:6:0:shutdown:/sbin:/sbin/shutdown
halt:x:7:0:halt:/sbin:/sbin/halt
mail:x:8:12:mail:/var/spool/mail:/sbin/nologin
operator:x:11:0:operator:/root:/sbin/nologin

[student@desktopX ~]$ tail -n 3 /etc/passwd
gdm:x:42:42:./var/lib/gdm:/sbin/nologin
gnome-initial-setup:x:993:991:./run/gnome-initial-setup:/sbin/nologin
tcpdump:x:72:72:./:/sbin/nologin
```

The **wc** command counts lines, words, and characters in a file. It can take a **-l**, **-w**, or **-c** option to display only the lines, words, or characters, respectively.

```
[student@desktopX ~]$ wc /etc/passwd
 39  70 2005 /etc/passwd

[student@desktopX ~]$ wc -l /etc/passwd
```

```
39 /etc/passwd
```

```
[student@desktopX ~]$ wc -c /etc/group /etc/hosts
```

```
843 /etc/group
```

```
227 /etc/hosts
```

```
1070 total
```

Tab completion

Tab completion allows a user to quickly complete commands or file names once they have typed enough at the prompt to make it unique. If the characters typed are not unique, pressing the **Tab** key twice displays all commands that begin with the characters already typed.

```
[student@desktopX ~]$ pas<Tab><Tab>
```

```
passwd      paste      pasuspender
```

```
[student@desktopX ~]$ pass<Tab>
```

```
[student@desktopX ~]$ passwd
```

```
Changing password for user student.
```

```
Changing password for student.
```

```
(current) UNIX password:
```

Tab completion can be used to complete file names when typing them as arguments to commands. When **Tab** is pressed, it will complete the file name as much as it can. Pressing **Tab** a second time causes the shell to list all of the files that are matched by the current pattern. Type additional characters until the name is unique, then use tab completion to finish off the command line.

```
[student@desktopX ~]$ ls /etc/pas<Tab>
```

```
[student@desktopX ~]$ ls /etc/passwd<Tab>
```

```
passwd  passwd-
```

Arguments and options can be matched with tab completion for many commands.

The **useradd** command is used by the superuser, **root**, to create additional users on the system. It has many options that can be used to control how that command behaves. Tab completion following a partial option can be used to complete the option without a lot of typing.

```
[root@desktopX ~]# useradd --<Tab><Tab>
```

```
--base-dir      --groups        --no-log-init   --shell
```

```
--comment       --help          --non-unique    --skel
```

```
--create-home   --home-dir      --no-user-group --system
```

```
--defaults      --inactive      --password      --uid
```

```
--expiredate    --key           --root          --user-group
```

```
--gid           --no-create-home --selinux-user
```

```
[root@desktopX ~]# useradd --
```


Command history

The **history** command displays a list of previously executed commands prefixed with a command number.

The exclamation point character, **!**, is a metacharacter that is used to expand previous commands without having to retype them. **!number** expands to the command matching the number specified. **!string** expands to the most recent command that begins with the string specified.

```
[student@desktopX ~]$ history
...Output omitted...
 23 clear
 24 who
 25 pwd
 26 ls /etc
 27 uptime
 28 ls -l
 29 date
 30 history

[student@desktopX ~]$ !ls
ls -l
total 0
drwxr-xr-x. 2 student student 6 Mar 29 21:16 Desktop
...Output omitted...

[student@desktopX ~]$ !26
ls /etc
abrt                hosts                pulse
adjtime             hosts.allow          purple
aliases             hosts.deny            qemu-ga
...Output omitted...
```

The arrow keys can be used to navigate through previous command lines in the shell's history. **Up Arrow** edits the previous command in the history list. **Down Arrow** edits the next command in the history list. Use this key when the **Up Arrow** has been pressed too many times. **Left Arrow** and **Right Arrow** move the cursor left and right in the current command line being edited.

The **Esc+.** key combination causes the shell to copy the last word of the previous command on the current command line where the cursor is. If used repeatedly, it will continue to go through earlier commands.

Editing the command line

When used interactively, **bash** has a command line-editing feature. This allows the user to use text editor commands to move around within and modify the current command being typed. Using the arrow

keys to move within the current command and to step through the command history was introduced earlier in this session. More powerful editing commands are introduced in the following table.

□

Useful command line-editing shortcuts

Shortcut	Description
Ctrl+a	Jump to the beginning of the command line.
Ctrl+e	Jump to the end of the command line.
Ctrl+u	Clear from the cursor to the beginning of the command line.
Ctrl+k	Clear from the cursor to the end of the command line.
Ctrl+Left Arrow	Jump to the beginning of the previous word on the command line.
Ctrl+Right Arrow	Jump to the beginning of the next word on the command line.
Ctrl+r	Search the history list of commands for a pattern.

There are several other command line-editing commands available, but these are the most useful commands for beginning users. The other commands can be found in the **bash(1)** man page.

References

bash(1), **date(1)**, **file(1)**, **head(1)**, **passwd(1)**, **tail(1)**, and **wc(1)** man pages

□ 1.6. Practice: Bash Commands and Keyboard Shortcuts

Quiz

Match the following Bash shortcuts to their descriptions in the table.

Description	Shell command
Jump to the beginning of the previous word on the command line.	Ctrl+Left Arrow
Separate commands on the same line.	;
Clear from the cursor to the end of the command line.	Ctrl+k
Re-execute a recent command by matching the command name.	!<i>string</i>
Shortcut used to complete commands, file names, and options.	Tab
Re-execute a specific command in the history list.	!<i>number</i>
Jump to the beginning for the command line.	Ctrl+a

Description	Shell command
Display the list of previous commands.	history
Copy the last argument of previous commands.	Esc+.

1.7. Lab: Accessing the Command Line

Performance checklist

In this lab, you will use the Bash shell to efficiently execute commands using shell metacharacters.

Resources:

Files:	/usr/bin/clean-binary-files
--------	------------------------------------

Outcomes:

- Practice using shell command line editing and history functions to efficiently execute commands with minor changes.
- Change the password of the **student** user to **T3st1ngT1me**.
- Execute commands used to identify file types and display parts of text files.

Before you begin...

Perform the following steps on desktopX.

1. Log into your **desktopX** system's graphical login screen as **student**.
2. Open a terminal window that will provide a **bash** prompt.
Select **Applications** → **Utilities** → **Terminal**.
3. Change **student**'s password to **T3st1ngT1me**.

Use the **passwd** command to change the password. Be sure to provide the original password, **student**, first.

```
[student@desktopX ~]$ passwd
Changing password for user student.
Changing password for student.
(current) UNIX password: student
New password: T3st1ngT1me
Retype new password: T3st1ngT1me
passwd: all authentication tokens updated successfully.
```

4. Display the current time and date.

```
5. [student@desktopX ~]$ date
6. Thu Apr 3 10:13:04 PDT 2014
```

7. Display the current time in the following format: HH:MM:SS A/PM. Hint: The format string that displays that output is **%r**.

Specify the **+%r** argument to **date**.

```
[student@desktopX ~]$ date +%r
10:14:07 AM
```

8. What kind of file is **/usr/bin/clean-binary-files**? Is it readable by humans?

Use the **file** command to determine its file type.

```
[student@desktopX ~]$ file
/usr/bin/clean-binary-files
/usr/bin/clean-binary-files: POSIX shell script, ASCII text executable
```

9. Use the **wc** command and **bash** shortcuts to display the size of **/usr/bin/clean-binary-files**.

The easiest shortcut to use is **Esc+.** to reuse the argument from the previous command.

```
[student@desktopX ~]$ wc <Esc>.
[student@desktopX ~]$ wc /usr/bin/clean-binary-files
 594 1780 13220 /usr/bin/clean-binary-files
```

10. Display the first 10 lines of **/usr/bin/clean-binary-files**.

The **head** command displays the beginning of the file. Did you use the **bash** shortcut again?

```
[student@desktopX ~]$ head <Esc>.
[student@desktopX ~]$ head
/usr/bin/clean-binary-files
#!/bin/sh
#
# Script to clean binary files.
#
# JPackage Project <http://www.jpackage.org/>
#
# $Id: clean-binary-files,v 1.1 2006/09/19 19:39:37 fnasser Exp $
#
# Import java functions
[ -r "/usr/share/java-utils/java-functions" ] \
```

11. Display the last 10 lines at the bottom of the **/usr/bin/clean-binary-files** file.

Use the **tail** command.

```
[student@desktopX ~]$ tail <Esc>.
[student@desktopX ~]$ tail
/usr/bin/clean-binary-files
...Output omitted...
```

12. Repeat the previous command, but use the **-n 20** option to display the last 20 lines in the file. Use command line editing to accomplish this with a minimal amount of keystrokes.

Up Arrow displays the previous command. **Ctrl+a** makes the cursor jump to the beginning of the line. **Ctrl+Right Arrow** jumps to the next word, then add the **-n 20** option and hit **Enter** to execute the command.

```
[student@desktopX ~]$ tail -n 20
/usr/bin/clean-binary-files
...Output omitted...
```

13. Execute the **date** command without any arguments to display the current date and time.

```
14. [student@desktopX ~]$ date
15. Thu Apr 3 10:48:30 PDT 2014
```

16. Use **bash** history to display just the time.

Display the list of previous commands with the **history** command to identify the specific **date** command to be executed. Execute the command with the **!number** history command.

```
[student@desktopX ~]$ history
...
 44 date +%X
...
[student@desktopX ~]$ !44
date +%X
10:49:56 AM
```

17. Finish your session with the **bash** shell.

Use either **exit** or the **Ctrl+d** key combination to close the shell.

```
[student@desktopX ~]$ exit
```

Chapter 2. Managing Files From the Command Line

2.1. The Linux File System Hierarchy

2.2. Practice: File System Hierarchy

2.3. Locating Files by Name

2.4. Practice: Locating Files and Directories

2.5. Managing Files Using Command-Line Tools

2.6. Practice: Command-Line File Management

2.7. Matching File Names Using Path Name Expansion

2.8. Practice: Path Name Expansion

2.9. Lab: Managing Files with Shell Expansion

Introduction

Overview	
Goal	To copy, move, create, delete, and organize files while working from the Bash shell prompt.
Objectives	<ul style="list-style-type: none">○ Identify the purpose for important directories on a Linux system.○ Specify files using absolute and relative path names.○ Create, copy, move, and remove files and directories using command-line utilities.○ Match one or more file names using shell expansion as arguments to shell commands.
Sections	<ul style="list-style-type: none">○ The Linux File System Hierarchy (and Practice)○ Locating Files by Name (and Practice)○ Managing Files Using Command-Line Tools (and Practice)○ Matching File Names Using Path Name Expansion (and Practice)
Lab	<ul style="list-style-type: none">○ Managing Files with Shell Expansion

2.1. The Linux File System Hierarchy

Objectives

After completing this section, students should be able to understand fundamental file system layout, organization, and the location of key file types.

The file system hierarchy

The Linux file system hierarchy

All files on a Linux system are stored on file systems which are organized into a single *inverted* tree of directories, known as a *file system hierarchy*. This tree is inverted because the **root of the tree** is said to be at the *top* of the hierarchy, and the branches of directories and subdirectories stretch *below* the root.

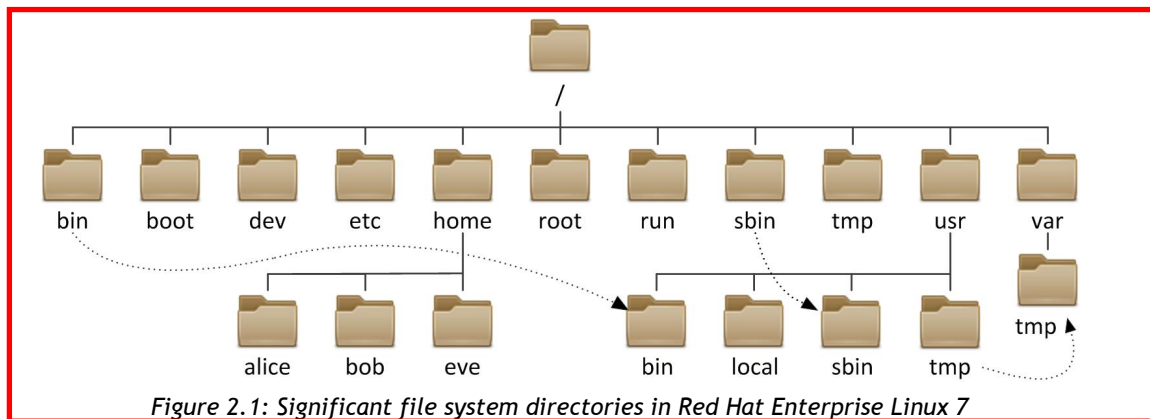


Figure 2.1: Significant file system directories in Red Hat Enterprise Linux 7

The directory **/** is the root directory at the top of the file system hierarchy. The **/** character is also used as a *directory separator* in file names. For example, if **etc** is a subdirectory of the **/** directory, we could call that directory **/etc**. Likewise, if the **/etc** directory contained a file named **issue**, we could refer to that file as **/etc/issue**.

Subdirectories of **/** are used for standardized purposes to organize files by type and purpose. This makes it easier to find files. For example, in the root directory, the subdirectory **/boot** is used for storing files needed to boot the system.

Note

The following terms are encountered in describing file system directory contents:

- *static* is content that remains unchanged until explicitly edited or reconfigured.
- *dynamic* or *variable* is content typically modified or appended by active processes.
- *persistent* is content, particularly configuration settings, that remain after a reboot.
- *runtime* is process- or system-specific content or attributes cleared during reboot.

The following table lists some of the most important directories on the system by name and purpose.

□

Important Red Hat Enterprise Linux directories

Location	Purpose
----------	---------

Location	Purpose
/usr	Installed software, shared libraries, include files, and static read-only program data. Important subdirectories include: <ul style="list-style-type: none"> - /usr/bin: <i>User commands.</i> - /usr/sbin: <i>System administration commands.</i> - /usr/local: <i>Locally customized software.</i>
/etc	Configuration files specific to this system.
/var	Variable data specific to this system that should persist between boots. Files that dynamically change (e.g. databases, cache directories, log files, printer-spoiled documents, and website content) may be found under /var .
/run	Runtime data for processes started since the last boot. This includes process ID files and lock files, among other things. The contents of this directory are recreated on reboot. (This directory consolidates /var/run and /var/lock from older versions of Red Hat Enterprise Linux.)
/home	Home directories where regular users store their personal data and configuration files.
/root	Home directory for the administrative superuser, root.
/tmp	A world-writable space for temporary files. Files which are more than 10 days old are deleted from this directory automatically. Another temporary directory exists, /var/tmp , in which files that have not been accessed, changed, or modified in more than 30 days are deleted automatically.
/boot	Files needed in order to start the boot process.
/dev	Contains special device files which are used by the system to access hardware.

Important

In Red Hat Enterprise Linux 7, four older directories in **/** now have identical contents as their counterparts located in **/usr**:

- **/bin** and **/usr/bin**.
- **/sbin** and **/usr/sbin**.

- **/lib** and **/usr/lib**.
- **/lib64** and **/usr/lib64**.

In older versions of Red Hat Enterprise Linux, these were distinct directories containing different sets of files. In RHEL 7, the directories in / are symbolic links to the matching directories in **/usr**.

References

hier(7) man page

Filesystem Hierarchy Standard

Quiz

Match the following items to their counterparts in the table.

Directory purpose	Location
This directory contains static, persistent system configuration data.	/etc
This is the system's root directory.	/
User home directories are located under this directory.	/home
This is the root account's home directory.	/root
This directory contains dynamic configuration data, such as FTP and websites.	/var
Regular user commands and utilities are located here.	/usr/bin
System administration binaries, for root use, are here.	/usr/sbin
Temporary files are stored here.	/tmp
Contains dynamic, non-persistent application runtime data.	/run
Contains installed software programs and libraries.	/usr

2.3. Locating Files by Name

Objectives

After completing this section, students should be able to correctly use absolute path names, change a working directory, and use commands to determine directory locations and contents.

Absolute paths and relative paths

Locating files by name



Figure 2.2: The common file browser view (left) is equivalent to the top-down view (right)

The *path* of a file or directory specifies its unique file system location. Following a file path traverses one or more named subdirectories, delimited by a forward slash (/), until the destination is reached. Standard file behavior definitions apply to directories (also called *folders*) the same as other file types.

Important

Although a **space** is an acceptable character in Linux file names, a space is the delimiter used by the command shell for command syntax interpretation. New administrators are advised to avoid using spaces in file names, since file names that include spaces frequently result in undesired command execution behavior.

Absolute paths

An *absolute path* is a *fully qualified* name, beginning at the root (/) directory and specifying each subdirectory traversed to reach and uniquely represent a single file. Every file in a file system has a unique absolute path name, recognized with a simple rule: A path name with a forward slash (/) as the first character is an absolute path name. For example, the absolute path name for the system message log file is **/var/log/messages**. Absolute path names can be long to type, so files may also be located *relatively*.

When a user logs in and opens a command window, the initial location is normally the user's home directory. System processes also have an initial directory. Users and processes navigate to other directories as needed; the terms *working directory* or *current working directory* refer to their *current* location.

Relative paths

Like an absolute path, a *relative path* identifies a unique file, specifying only the path necessary to reach the file from the working directory. Recognizing relative path names follows a simple rule: A path name with *anything other than* a forward slash (/) as a first character is a relative path name. A user in the **/var** directory could refer to the message log file relatively as **log/messages**.

For standard Linux file systems, the path name of a file, including all / characters, may be no more than 4095 bytes long. Each component of the path name separated by / characters may be no more than 255 bytes long. File names can use any UTF-8 encoded Unicode character except / and the **NUL** character. (ASCII characters require one byte; other Latin, Greek, Hebrew, or Cyrillic characters take two bytes; remaining characters in the Unicode Basic Multilingual Plane take three; and no character will take more than four bytes.)

Linux file systems—including, but not limited to, ext4, XFS, BTRFS, GFS2, and GlusterFS—are case-sensitive. Creating `FileCase.txt` and `filecase.txt` in the same directory results in two unique files. Although many non-Linux file systems are supported in Linux, each has unique file naming rules. For example, the ubiquitous VFAT file system is not case-sensitive and allows only one of the two example files to be created. However, VFAT, along with Microsoft's NTFS and Apple's HFS+, has *case preserving* behavior. Although these file systems are *not* case-sensitive (enforced primarily to support backward compatibility), they do display file names with the original capitalization used when the file was created.

▮ Navigating paths

The `pwd` command displays the full path name of the current location, which helps determine appropriate syntax for reaching files using relative path names. The `ls` command lists directory contents for the specified directory or, if no directory is given, for the current directory.

```
[student@desktopX ~]$ pwd
/home/student

[student@desktopX ~]$ ls
Desktop Documents Downloads Music Pictures Public Templates Videos

[student@desktopX ~]$
```

Use the `cd` command to change directories. With a working directory of `/home/student`, relative path syntax is shortest to reach the `Videos` subdirectory. The `Documents` subdirectory is then reached using absolute path syntax.

```
[student@desktopX ~]$ cd Videos
[student@desktopX Videos]$ pwd
/home/student/Videos

[student@desktopX Videos]$ cd /home/student/Documents
[student@desktopX Documents]$ pwd
/home/student/Documents

[student@desktopX Documents]$ cd
[student@desktopX ~]$ pwd
/home/student

[student@desktopX ~]$
```

The shell program prompt displays, for brevity, only the last component of the current directory path. For `/home/student/Videos`, only `Videos` displays. At any time, return to the user's home directory using `cd` without specifying a destination. The prompt displays the *tilde* (`~`) character when the user's current directory is their home directory.

The `touch` command normally updates a file's timestamp to the current date and time without otherwise modifying it. This is useful for creating empty files, which can be used for practice, since "touching" a file name that does not exist causes the file to be created. Using `touch`, practice files are created in the `Documents` and `Videos` subdirectories.

```
[student@desktopX ~]$ touch Videos/blockbuster1.ogg
[student@desktopX ~]$ touch Videos/blockbuster2.ogg
[student@desktopX ~]$ touch Documents/thesis_chapter1.odf
[student@desktopX ~]$ touch Documents/thesis_chapter2.odf
[student@desktopX ~]$
```

The **ls** command has multiple options for displaying attributes on files. The most common and useful are **-l** (long listing format), **-a** (all files, includes *hidden* files), and **-R** (recursive, to include the contents of all subdirectories).

```
[student@desktopX ~]$ ls -l
total 15
drwxr-xr-x. 2 student student 4096 Feb  7 14:02 Desktop
drwxr-xr-x. 2 student student 4096 Jan  9 15:00 Documents
drwxr-xr-x. 3 student student 4096 Jan  9 15:00 Downloads
drwxr-xr-x. 2 student student 4096 Jan  9 15:00 Music
drwxr-xr-x. 2 student student 4096 Jan  9 15:00 Pictures
drwxr-xr-x. 2 student student 4096 Jan  9 15:00 Public
drwxr-xr-x. 2 student student 4096 Jan  9 15:00 Templates
drwxr-xr-x. 2 student student 4096 Jan  9 15:00 Videos

[student@desktopX ~]$ ls -a
total 15
drwx-----. 16 student student  4096 Feb  8 16:15 .
drwxr-xr-x.  6 root    root      4096 Feb  8 16:13 ..
-rw-----.  1 student student 22664 Feb  8 00:37 .bash_history
-rw-r--r--.  1 student student   18 Jul  9  2013 .bash_logout
-rw-r--r--.  1 student student  176 Jul  9  2013 .bash_profile
-rw-r--r--.  1 student student  124 Jul  9  2013 .bashrc
drwxr-xr-x.  4 student student  4096 Jan 20 14:02 .cache
drwxr-xr-x.  8 student student  4096 Feb  5 11:45 .config
drwxr-xr-x.  2 student student  4096 Feb  7 14:02 Desktop
drwxr-xr-x.  2 student student  4096 Jan  9 15:00 Documents
drwxr-xr-x.  3 student student  4096 Jan 25 20:48 Downloads
drwxr-xr-x. 11 student student  4096 Feb  6 13:07 .gnome2
drwx-----.  2 student student  4096 Jan 20 14:02 .gnome2_private
-rw-----.  1 student student 15190 Feb  8 09:49 .ICEauthority
drwxr-xr-x.  3 student student  4096 Jan  9 15:00 .local
drwxr-xr-x.  2 student student  4096 Jan  9 15:00 Music
```

```
drwxr-xr-x. 2 student student 4096 Jan 9 15:00 Pictures
drwxr-xr-x. 2 student student 4096 Jan 9 15:00 Public
drwxr-xr-x. 2 student student 4096 Jan 9 15:00 Templates
drwxr-xr-x. 2 student student 4096 Jan 9 15:00 Videos
```

```
[student@desktopX ~]$
```

The two special directories at the top of the listing refer to the current directory (`.`) and the *parent* directory (`..`). These special directories exist in every directory on the system. Their usefulness will become apparent when file management commands are practiced.

Important

File names beginning with a dot (`.`) indicate files *hidden* from normal view using `ls` and other commands. This *is not* a security feature. Hidden files keep necessary user configuration files from cluttering home directories. Many commands process hidden files only with specific command-line options, preventing one user's configuration from being accidentally copied to other directories or users.

To protect file *contents* from improper viewing requires the use of *file permissions*.

```
[student@desktopX ~]$ ls -R
```

```
..:
```

```
Desktop Documents Downloads Music Pictures Public Templates Videos
```

```
./Desktop:
```

```
./Documents:
```

```
thesis_chapter1.odf thesis_chapter2.odf
```

```
./Downloads:
```

```
./Music:
```

```
./Pictures:
```

```
./Public:
```

```
./Templates:
```

```
./Videos:
```

```
blockbuster1.ogg blockbuster2.ogg
```

```
[student@desktopX ~]$
```

The **cd** command has many options. A few are so useful as to be worth practicing early and using often. The command **cd -** changes directory to the directory where the user was *previous* to the current directory. Watch as this user takes advantage of this behavior to alternate between two directories, useful when processing a series of similar tasks.

```
[student@desktopX ~]$ cd Videos
[student@desktopX Videos]$ pwd
/home/student/Videos
[student@desktopX Videos]$ cd /home/student/Documents
[student@desktopX Documents]$ pwd
/home/student/Documents
[student@desktopX Documents]$ cd -
[student@desktopX Videos]$ pwd
/home/student/Videos
[student@desktopX Videos]$ cd -
[student@desktopX Documents]$ pwd
/home/student/Documents
[student@desktopX Documents]$ cd -
[student@desktopX Videos]$ pwd
/home/student/Videos
[student@desktopX Videos]$ cd
[student@desktopX ~]$
```

The **cd ..** command uses the **..** hidden directory to move up one level to the *parent* directory, without needing to know the exact parent name. The other hidden directory (**.**) specifies the *current directory* on commands in which the current location is either the source or destination argument, avoiding the need to type out the directory's absolute path name.

```
[student@desktopX Videos]$ pwd
/home/student/Videos
[student@desktopX Videos]$ cd .
[student@desktopX Videos]$ pwd
/home/student/Videos
[student@desktopX Videos]$ cd ..
[student@desktopX ~]$ pwd
/home/student
[student@desktopX ~]$ cd ..
[student@desktopX home]$ pwd
/home
[student@desktopX home]$ cd ..
```

```
[student@desktopX ~]$ pwd
/
[student@desktopX ~]$ cd
[student@desktopX ~]$ pwd
/home/student
[student@desktopX ~]$
```

References

info libc 'file name resolution' (*GNU C Library Reference Manual*)

- o Section 11.2.2 File name resolution

bash(1), cd(1), ls(1), pwd(1), unicode(7), and utf-8(7) man pages

[UTF-8 and Unicode](#)

Quiz

Match the following items to their counterparts in the table.

Action to accomplish	Command
List the current user's home directory (long format) in simplest syntax, when it is not the current location.	ls -l ~
Return to the current user's home directory.	cd
Determine the absolute path name of the current location.	pwd
Return to the most previous working directory.	cd -
Move up two levels from the current location.	cd ../../
List the current location (long format) with hidden files.	ls -al
Move to the binaries location, from any current location.	cd /bin
Move up to the parent of the current location.	cd ..
Move to the binaries location, from the root directory.	cd bin

2.5. Managing Files Using Command-Line Tools

Objectives

After completing this section, students should be able to create, copy, link, move, and remove files and subdirectories in various directories.

Command-line file management

Managing files using command-line tools

File management involves creating, deleting, copying, and moving files. Additionally, directories can be created, deleted, copied, and moved to help organize files logically. When working at the command line, file management requires awareness of the current working directory to choose either absolute or relative path syntax as most efficient for the immediate task.

□

File management commands

Activity	Single source ^(note)	Multiple source ^(note)
Copy file	cp file1 file2	cp file1 file2 file3 dir ⁽⁵⁾
Move file	mv file1 file2 ⁽¹⁾	mv file1 file2 file3 dir ⁽⁴⁾
Remove file	rm file1	rm -f file1 file2 file3 ⁽⁵⁾
Create directory	mkdir dir	mkdir -p par1/par2/dir ⁽⁶⁾
Copy directory	cp -r dir1 dir2 ⁽²⁾	cp -r dir1 dir2 dir3 dir4 ⁽⁴⁾
Move directory	mv dir1 dir2 ⁽³⁾	mv dir1 dir2 dir3 dir4 ⁽⁴⁾
Remove directory	rm -r dir1 ⁽²⁾	rm -rf dir1 dir2 dir3 ⁽⁵⁾
Note:	<p>⁽¹⁾The result is a rename.</p> <p>⁽²⁾The "recursive" option is required to process a source directory.</p> <p>⁽³⁾If dir2 exists, the result is a move. If dir2 doesn't exist, the result is a rename.</p> <p>⁽⁴⁾The last argument must be a directory.</p> <p>⁽⁵⁾Use caution with "force" option; you will not be prompted to confirm your action.</p> <p>⁽⁶⁾Use caution with "create parent" option; typing errors are not caught.</p>	

Create directories

The **mkdir** command creates one or more directories or subdirectories, generating errors if the file name already exists or when attempting to create a directory in a parent directory that doesn't exist. The **-p** *parent* option creates missing parent directories for the requested destination. Be cautious when using **mkdir -p**, since accidental spelling mistakes create unintended directories without generating error messages.

```
[student@desktopX ~]$ mkdir Video/Watched
```

```
mkdir: cannot create directory `Video/Watched': No such file or directory
```

The **mkdir** failed because **Videos** was misspelled. "**Video**" does not exist as a location in which to create the **Watched** subdirectory. If a **-p** were used, the user would not have received an error message and now have two directories, **Video** and **Videos**.

```
[student@desktopX ~]$ mkdir Videos/Watched
```

```
[student@desktopX ~]$ cd Documents
```



```

[student@desktopX Documents]$ mkdir ProjectX ProjectY
[student@desktopX Documents]$ mkdir -p Thesis/Chapter1
Thesis/Chapter2 Thesis/Chapter3
[student@desktopX Documents]$ cd
[student@desktopX ~]$ ls -R Videos Documents
Documents:
ProjectX ProjectY Thesis thesis_chapter1.odf thesis_chapter2.odf

Documents/ProjectX:

Documents/ProjectY:

Documents/Thesis:
Chapter1 Chapter2 Chapter3

Documents/Thesis/Chapter1:

Documents/Thesis/Chapter2:

Documents/Thesis/Chapter3:

Videos:
blockbuster1.ogg blockbuster2.ogg Watched

Videos/Watched:

[student@desktopX ~]$

```

The last **mkdir** created three Chapter_N subdirectories with one command. The **-p** *parent* option created the missing parent directory **Thesis**.

Copy files

The **cp** command copies one or more files to become new, independent files. Syntax allows copying an existing file to a new file in the current or another directory, or copying multiple files into another directory. In any destination, new file names must be unique. If the new file name is not unique, the copy command will overwrite the existing file.

```

[student@desktopX ~]$ cd Videos
[student@desktopX Videos]$ cp blockbuster1.ogg
blockbuster3.ogg

```

```

[student@desktopX Videos]$ ls -l
total 0
-rw-rw-r--. 1 student student 0 Feb 8 16:23 blockbuster1.ogg
-rw-rw-r--. 1 student student 0 Feb 8 16:24 blockbuster2.ogg
-rw-rw-r--. 1 student student 0 Feb 8 19:02 blockbuster3.ogg
drwxrwxr-x. 2 student student 4096 Feb 8 23:35 Watched
[student@desktopX Videos]$

```

When copying multiple files with one command, the last argument must be a directory. Copied files retain their original names in the new directory. Conflicting file names that exist at a destination may be overwritten. To protect users from accidentally overwriting directories with contents, multiple file **cp** commands ignore directories specified as a source. Copying non-empty directories, with contents, requires the **-r** *recursive* option.

```

[student@desktopX Videos]$ cd ../Documents
[student@desktopX Documents]$ cp thesis_chapter1.odf
thesis_chapter2.odf Thesis ProjectX
cp: omitting directory `Thesis'
[student@desktopX Documents]$ cp -r Thesis ProjectX
[student@desktopX Documents]$ cp thesis_chapter2.odf
Thesis/Chapter2/
[student@desktopX Documents]$ ls -R
.:
ProjectX ProjectY Thesis thesis_chapter1.odf thesis_chapter2.odf

./ProjectX:
Thesis thesis_chapter1.odf thesis_chapter2.odf

./ProjectX/Thesis:

./ProjectY:

./Thesis:
Chapter1 Chapter2 Chapter3

./Thesis/Chapter1:

./Thesis/Chapter2:

```

```
thesis_chapter2.odf
```

```
./Thesis/Chapter3:
```

```
[student@desktopX Documents]$
```

In the first **cp** command, **Thesis** failed to copy, but **thesis_chapter1.odf** and **thesis_chapter2.odf** succeeded. Using the **-r recursive** option, copying **Thesis** succeeded.

Move files

The **mv** command renames files in the same directory, or relocates files to a new directory. File contents remain unchanged. Files moved to a different file system require creating a new file by copying the source file, then deleting the source file. Although normally transparent to the user, large files may take noticeably longer to move.

```
[student@desktopX Videos]$ cd ../Documents
```

```
[student@desktopX Documents]$ ls -l
```

```
total 0
```

```
-rw-rw-r--. 1 student student 0 Feb 8 16:24 thesis_chapter1.odf
```

```
-rw-rw-r--. 1 student student 0 Feb 8 16:24 thesis_chapter2.odf
```

```
[student@desktopX Documents]$ mv thesis_chapter2.odf  
thesis_chapter2_reviewed.odf
```

```
[student@desktopX Documents]$ mv thesis_chapter1.odf  
Thesis/Chapter1
```

```
[student@desktopX Documents]$ ls -lR
```

```
..:
```

```
total 16
```

```
drwxrwxr-x. 2 student student 4096 Feb 11 11:58 ProjectX
```

```
drwxrwxr-x. 2 student student 4096 Feb 11 11:55 ProjectY
```

```
drwxrwxr-x. 5 student student 4096 Feb 11 11:56 Thesis
```

```
-rw-rw-r--. 1 student student 0 Feb 11 11:54 thesis_chapter2_reviewed.odf
```

```
./ProjectX:
```

```
total 0
```

```
-rw-rw-r--. 1 student student 0 Feb 11 11:58 thesis_chapter1.odf
```

```
-rw-rw-r--. 1 student student 0 Feb 11 11:58 thesis_chapter2.odf
```

```
./ProjectX/Thesis:
```

```
total 0
```

```

./ProjectY:
total 0

./Thesis:
total 12
drwxrwxr-x. 2 student student 4096 Feb 11 11:59 Chapter1
drwxrwxr-x. 2 student student 4096 Feb 11 11:56 Chapter2
drwxrwxr-x. 2 student student 4096 Feb 11 11:56 Chapter3

./Thesis/Chapter1:
total 0
-rw-rw-r--. 1 student student 0 Feb 11 11:54 thesis_chapter1.odf

./Thesis/Chapter2:
total 0
-rw-rw-r--. 1 student student 0 Feb 11 11:54 thesis_chapter2.odf

./Thesis/Chapter3:
total 0
[student@desktopX Documents]$

```

The first **mv** command is an example of renaming a file. The second causes the file to be relocated to another directory.

Remove files and directories

Default syntax for **rm** deletes files, but not directories. Deleting a directory, and potentially many subdirectories and files below it, requires the **-r recursive** option. There is no command-line undelete feature, nor a trash bin from which to restore.

```

[student@desktopX Documents]$ pwd
/home/student/Documents
[student@desktopX Documents]$ rm
thesis_chapter2_reviewed.odf
[student@desktopX Documents]$ rm Thesis/Chapter1
rm: cannot remove `Thesis/Chapter1': Is a directory
[student@desktopX Documents]$ rm -r Thesis/Chapter1
[student@desktopX Documents]$ ls -l Thesis
total 8
drwxrwxr-x. 2 student student 4096 Feb 11 12:47 Chapter2

```

```

drwxrwxr-x. 2 student student 4096 Feb 11 12:48 Chapter3
[student@desktopX Documents]$ rm -ri Thesis
rm: descend into directory `Thesis'? y
rm: descend into directory `Thesis/Chapter2'? y
rm: remove regular empty file `Thesis/Chapter2/thesis_chapter2.odf'? y
rm: remove directory `Thesis/Chapter2'? y
rm: remove directory `Thesis/Chapter3'? y
rm: remove directory `Thesis'? y
[student@desktopX Documents]$

```

After **rm** failed to delete the **Chapter1** directory, the **-r recursive** option succeeded. The last **rm** command parsed into each subdirectory first, individually deleting contained files before removing each now-empty directory. Using **-i** will interactively prompt for each deletion. This is essentially the opposite of **-f** which will force the deletion without prompting the user. The **rmdir** command deletes directories only if empty. Removed directories cannot be undeleted.

```

[student@desktopX Documents]$ pwd
/home/student/Documents
[student@desktopX Documents]$ rmdir ProjectY
[student@desktopX Documents]$ rmdir ProjectX
rmdir: failed to remove `ProjectX': Directory not empty
[student@desktopX Documents]$ rm -r ProjectX
[student@desktopX Documents]$ ls -lR
.:
total 0
[student@desktopX Documents]$

```

The **rmdir** command failed to delete non-empty **ProjectX**, but **rm -r** succeeded.

References

[cp\(1\)](#), [mkdir\(1\)](#), [mv\(1\)](#), [rm\(1\)](#), and [rmdir\(1\)](#) man pages

2.6. Practice: Command-Line File Management

Guided exercise

In this lab, you will practice efficient techniques for creating and organizing files using directories, file copies, and links.

Outcomes:

Students will practice creating, rearranging, and deleting files.

Before you begin...

Log into your student account on serverX. Begin in your home directory.

1. In your home directory, create sets of empty practice files to use for the remainder of this lab. If the intended command is not immediately recognized, students are expected to use the guided solution to see and practice how the task is accomplished. Use the shell tab completion to locate and complete path names more easily.

Create six files with names of the form **songX.mp3**.

Create six files with names of the form **snapX.jpg**.

Create six files with names of the form **filmX.avi**.

In each set, replace X with the numbers 1 through 6.

```
[student@serverX ~]$ touch song1.mp3 song2.mp3
song3.mp3 song4.mp3 song5.mp3 song6.mp3
[student@serverX ~]$ touch snap1.jpg snap2.jpg
snap3.jpg snap4.jpg snap5.jpg snap6.jpg
[student@serverX ~]$ touch film1.avi film2.avi
film3.avi film4.avi film5.avi film6.avi
[student@serverX ~]$ ls -l
```

2. From your home directory, move the song files into your **Music** subdirectory, the snapshot files into your **Pictures** subdirectory, and the movie files into your **Videos** subdirectory.

When distributing files from one location to many locations, first change to the directory containing the *source* files. Use the simplest path syntax, absolute or relative, to reach the destination for each file management task.

```
[student@serverX ~]$ mv song1.mp3 song2.mp3 song3.mp3
song4.mp3 song5.mp3 song6.mp3 Music
[student@serverX ~]$ mv snap1.jpg snap2.jpg snap3.jpg
snap4.jpg snap5.jpg snap6.jpg Pictures
[student@serverX ~]$ mv film1.avi film2.avi film3.avi
film4.avi film5.avi film6.avi Videos
[student@serverX ~]$ ls -l Music Pictures Videos
```

3. In your home directory, create three subdirectories for organizing your files into projects. Call these directories **friends**, **family**, and **work**. Create all three with one command.

You will use these directories to rearrange your files into projects.

```
[student@serverX ~]$ mkdir friends family work
[student@serverX ~]$ ls -l
```

4. You will collect some of the new files into the project directories for family and friends. Use as many commands as needed. You do not have to use only one command as in the example. For each project, first change to the project directory, then copy the source files into this

directory. You are making copies, since you will keep the originals after giving these projects to family and friends.

Copy files (all types) containing numbers 1 and 2 to the friends folder.

Copy files (all types) containing numbers 3 and 4 to the family folder.

When collecting files from multiple locations into one location, change to the directory that will contain the *destination* files. Use the simplest path syntax, absolute or relative, to reach the source for each file management task.

```
[student@serverX ~]$ cd friends
[student@serverX friends]$ cp ~/Music/song1.mp3
~/Music/song2.mp3 ~/Pictures/snap1.jpg
~/Pictures/snap2.jpg ~/Videos/film1.avi
~/Videos/film2.avi .
[student@serverX friends]$ ls -l
[student@serverX friends]$ cd ../family
[student@serverX family]$ cp ~/Music/song3.mp3
~/Music/song4.mp3 ~/Pictures/snap3.jpg
~/Pictures/snap4.jpg ~/Videos/film3.avi
~/Videos/film4.avi .
[student@serverX family]$ ls -l
```

5. For your work project, you will create additional copies.

```
6. [student@serverX family]$ cd ../work
7. [student@serverX work]$ cp ~/Music/song5.mp3
~/Music/song6.mp3 ~/Pictures/snap5.jpg
~/Pictures/snap6.jpg ~/Videos/film5.avi
~/Videos/film6.avi .
```

```
[student@serverX work]$ ls -l
```

8. Your projects are now done. Time to clean up the projects.

Change to your home directory. Attempt to delete both the family and friends projects with a single **rmdir** command.

```
[student@serverX work]$ cd
[student@serverX ~]$ rmdir family friends
rmdir: failed to remove `family': Directory not empty
rmdir: failed to remove `friends': Directory not empty
```

Using the **rmdir** command should fail since both directories are non-empty.

9. Use another command that will succeed in deleting both the family and friends folders.

```
10. [student@serverX ~]$ rm -r family friends
```

```
[student@serverX ~]$ ls -l
```

11. Delete all the files in the work project, but do not delete the work directory.

```
12. [student@serverX ~]$ cd work
```

```
13. [student@serverX work]$ rm song5.mp3 song6.mp3  
snap5.jpg snap6.jpg film5.avi film6.avi
```

```
[student@serverX work]$ ls -l
```

14. Finally, from your home directory, use the **rmdir** command to delete the work directory. The command should succeed now that it is empty.

```
15. [student@serverX work]$ cd
```

```
16. [student@serverX ~]$ rmdir work
```

```
[student@serverX ~]$ ls -l
```

2.7. Matching File Names Using Path Name Expansion

Objectives

After completing this section, students should be able to use meta-characters and expansion techniques to improve file management processing efficiency.

File globbing: path name expansion

Matching file names using path name expansion

The Bash shell has a path name-matching capability historically called *globbing*, abbreviated from the “global command” file path expansion program of early UNIX. The Bash globbing feature, commonly called *pattern matching* or “wildcards”, makes managing large numbers of files easier. Using *meta-characters* that “expand” to match file and path names being sought, commands perform on a focused set of files at once.

Pattern matching

Globbing is a shell command-parsing operation that expands a wildcard pattern into a list of matching path names. Command-line meta-characters are replaced by the match list prior to command execution. Patterns, especially square-bracketed character classes, that do not return matches display the original pattern request as literal text. The following are common meta-characters and pattern classes.

Pattern	Matches
*	Any string of 0 or more characters.
?	Any single character.
~	The current user’s home directory.

Pattern	Matches
<code>~username</code>	User <code>username</code> 's home directory.
<code>~+</code>	The current working directory.
<code>~-</code>	The previous working directory.
<code>[abc...]</code>	Any one character in the enclosed class.
<code>[!abc...]</code>	Any one character <i>not</i> in the enclosed class.
<code>[^abc...]</code>	Any one character <i>not</i> in the enclosed class.
<code>[:alpha:]</code>	Any alphabetic character. ⁽¹⁾ 大小写字母
<code>[:lower:]</code>	Any lower-case character. ⁽¹⁾
<code>[:upper:]</code>	Any upper-case character. ⁽¹⁾
<code>[:alnum:]</code>	Any alphabetic character or digit. ⁽¹⁾
<code>[:punct:]</code>	Any printable character not a space or alphanumeric. ⁽¹⁾ 标点符号
<code>[:digit:]</code>	Any digit, 0-9 . ⁽¹⁾
<code>[:space:]</code>	Any one whitespace character; may include tabs , newline, or carriage returns, and form feeds as well as space. ⁽¹⁾
Note	⁽¹⁾ pre-set POSIX character class; adjusts for current locale.

A sample set of files is useful to demonstrate expansion.

```
[student@desktopX ~]$ mkdir glob; cd glob
[student@desktopX glob]$ touch alfa bravo charlie delta echo
able baker cast dog easy
[student@desktopX glob]$ ls
able alfa baker bravo cast charlie delta dog easy echo
[student@desktopX glob]$
```

First, simple pattern matches using `*` and `?`.

```
[student@desktopX glob]$ ls a*
able alfa

[student@desktopX glob]$ ls *a*
able alfa baker bravo cast charlie delta easy

[student@desktopX glob]$ ls [ac]*
able alfa cast charlie

[student@desktopX glob]$ ls ????
```

```
able alfa cast easy echo
[student@desktopX glob]$ ls ?????
baker bravo delta
[student@desktopX glob]$
```

Tilde expansion

The tilde character (~), when followed by a slash delimiter, matches the current user's home directory. When followed by a string of characters up to a slash, it will be interpreted as a username, if one matches. If no username matches, then an actual tilde followed by the string of characters will be returned.

```
[student@desktopX glob]$ ls ~/glob
able alfa baker bravo cast charlie delta dog easy echo
[student@desktopX glob]$ echo ~/glob
/home/student/glob
[student@desktopX glob]$
```

Brace expansion

Brace expansion is used to generate discretionary strings of characters. Braces contain a comma-separated list of strings, or a sequence expression. The result includes the text preceding or following the brace definition. Brace expansions may be nested, one inside another.

```
[student@desktopX glob]$ echo
{Sunday,Monday,Tuesday,Wednesday}.log
Sunday.log Monday.log Tuesday.log Wednesday.log
[student@desktopX glob]$ echo file{1..3}.txt
file1.txt file2.txt file3.txt
[student@desktopX glob]$ echo file{a..c}.txt
filea.txt fileb.txt filec.txt
[student@desktopX glob]$ echo file{a,b}{1,2}.txt
filea1.txt filea2.txt fileb1.txt fileb2.txt
[student@desktopX glob]$ echo file{a{1,2},b,c}.txt
filea1.txt filea2.txt fileb.txt filec.txt
[student@desktopX glob]$
```

Command substitution

Command substitution allows the output of a command to replace the command itself. Command substitution occurs when a command is enclosed with a beginning dollar sign and parenthesis, **\$ (command)**, or with backticks, **`command`**. The form with backticks is older, and has two disadvantages: 1) it can be easy to visually confuse backticks with single quote marks, and 2) backticks cannot be nested inside backticks. The **\$ (command)** form can nest multiple command expansions inside each other.

```
[student@desktopX glob]$ echo Today is `date +%A`.
```

```
Today is Wednesday.
```

```
[student@desktopX glob]$ echo The time is $(date +%M) minutes  
past $(date +%l%p) .
```

```
The time is 26 minutes past 11AM.
```

```
[student@desktopX glob]$
```

Protecting arguments from expansion

Many characters have special meaning in the Bash shell. To ignore meta-character special meanings, *quoting* and *escaping* are used to protect them from shell expansion. The backslash (\) is an escape character in Bash, protecting the single following character from special interpretation. To protect longer character strings, single (') or double quotes (") are used to enclose strings. Use double quotation marks to suppress globbing and shell expansion, but still allow command and variable substitution. Variable substitution is conceptually identical to command substitution, but may use optional brace syntax.

```
[student@desktopX glob]$ host=$(hostname); echo $host
```

```
desktopX
```

```
[student@desktopX glob]$ echo "***** hostname is ${host}  
*****"
```

```
***** hostname is desktopX *****
```

```
[student@desktopX glob]$ echo Your username variable is  
\$USER.
```

```
Your username variable is $USER.
```

```
[student@desktopX glob]$
```

Use single quotation marks to interpret *all* text literally. Observe the difference, on both screen and keyboard, between the single quote (') and the command substitution backtick (`). Besides suppressing globbing and shell expansion, quotations direct the shell to additionally suppress command and variable substitution. The question mark is a meta-character that also needed protection from expansion.

```
[student@desktopX glob]$ echo "Will variable $host evaluate  
to $(hostname)?"
```

```
Will variable desktopX evaluate to desktopX?
```

```
[student@desktopX glob]$ echo 'Will variable $host evaluate  
to $(hostname)?'
```

```
Will variable $host evaluate to $(hostname)?
```

```
[student@desktopX glob]$
```

References

bash(1), **cd(1)**, **glob(7)**, **isalpha(3)**, **ls(1)**, **path_resolution(7)**, and **pwd(1)**
man pages

2.8. Practice: Path Name Expansion

Quiz

Match the following items to their counterparts in the table.

Requested match to find	Patterns
Only filenames beginning with "b"	b*
Only filenames ending in "b"	*b
Only filenames containing a "b"	*b*
Only filenames where first character is not "b"	[!b]*
Only filenames at least 3 characters in length	???*
Only filenames that contain a number	*[[:digit:]]*
Only filenames that begin with an upper-case letter	[[:upper:]]*

2.9. Lab: Managing Files with Shell Expansion

Performance checklist

In this lab, you will create, move, and remove files and folders using a variety of file name matching shortcuts.

Outcomes:

Familiarity and practice with many forms of wildcards for locating and using files.

Before you begin...

Perform the following steps on serverX unless directed otherwise. Log in as **student** and begin the lab in the home directory.

1. To begin, create sets of empty practice files to use in this lab. If an intended shell expansion shortcut is not immediately recognized, students are expected to use the solution to learn and practice. Use shell tab completion to locate file path names easily.

Create a total of 12 files with names **tv_seasonX_episodeY.ogg**. Replace X with the season number and Y with that season's episode, for two seasons of six episodes each.

```
[student@serverX ~]$ touch
tv_season{1..2}_episode{1..6}.ogg
[student@serverX ~]$ ls -l
```

2. As the author of a successful series of mystery novels, your next bestseller's chapters are being edited for publishing. Create a total of eight files with names **mystery_chapterX.odf**. Replace X with the numbers 1 through 8.

```
3. [student@serverX ~]$ touch mystery_chapter{1..8}.odf
```

```
[student@serverX ~]$ ls -l
```

- To organize the TV episodes, create two subdirectories named **season1** and **season2** under the existing **Videos** directory. Use one command.

```
5. [student@serverX ~]$ mkdir Videos/season{1..2}
```

```
[student@serverX ~]$ ls -lR
```

- Move the appropriate TV episodes into the season subdirectories. Use only two commands, specifying destinations using relative syntax.

```
7. [student@serverX ~]$ mv tv_season1* Videos/season1
```

```
8. [student@serverX ~]$ mv tv_season2* Videos/season2
```

```
[student@serverX ~]$ ls -lR
```

- To organize the mystery book chapters, create a two-level directory hierarchy with one command. Create **my_bestseller** under the existing **Documents** directory, and **chapters** beneath the new **my_bestseller** directory.

```
10. [student@serverX ~]$ mkdir -p  
Documents/my_bestseller/chapters
```

```
[student@serverX ~]$ ls -lR
```

- Using one command, create three more subdirectories directly under the **my_bestseller** directory. Name these subdirectories **editor**, **plot_change**, and **vacation**. The *create parent* option is not needed since the **my_bestseller** parent directory already exists.

```
12. [student@serverX ~]$ mkdir  
Documents/my_bestseller/{editor,plot_change,vacatio  
n}
```

```
[student@serverX ~]$ ls -lR
```

- Change to the **chapters** directory. Using the home directory shortcut to specify the source files, move all book chapters into the **chapters** directory, which is now your current directory. What is the simplest syntax to specify the destination directory?

```
14. [student@serverX ~]$ cd  
Documents/my_bestseller/chapters
```

```
15. [student@serverX chapters]$ mv ~/mystery_chapter* .
```

```
[student@serverX chapters]$ ls -l
```

16. The first two chapters are sent to the editor for review. To remember to not modify these chapters during the review, move those two chapters only to the **editor** directory. Use relative syntax starting from the **chapters** subdirectory.

```
17. [student@serverX chapters]$ mv mystery_chapter1.odf
    mystery_chapter2.odf ../editor
18. [student@serverX chapters]$ ls -l
```

```
[student@serverX chapters]$ ls -l ../editor
```

19. Chapters 7 and 8 will be written while on vacation. Move the files from **chapters** to **vacation**. Use one command without wildcard characters.

```
20. [student@serverX chapters]$ mv mystery_chapter7.odf
    mystery_chapter8.odf ../vacation
21. [student@serverX chapters]$ ls -l
```

```
[student@serverX chapters]$ ls -l ../vacation
```

22. With one command, change the working directory to the season 2 TV episodes location, then copy the first episode of the season to the **vacation** directory.

```
23. [student@serverX chapters]$ cd ~/Videos/season2
```

```
[student@serverX season2]$ cp tv_season2_episode1.ogg
~/Documents/my_bestseller/vacation
```

24. With one command, change the working directory to **vacation**, then list its files. Episode 2 is also needed. Return to the **season2** directory using the *previous working directory* shortcut. This will succeed if the last directory change was accomplished with one command. Copy the episode 2 file into **vacation**. Return to **vacation** using the shortcut again.

```
25. [student@serverX season2]$ cd
    ~/Documents/my_bestseller/vacation
26. [student@serverX vacation]$ ls -l
27. [student@serverX vacation]$ cd -
28. [student@serverX season2]$ cp tv_season2_episode2.ogg
    ~/Documents/my_bestseller/vacation
29. [student@serverX vacation]$ cd -
```

```
[student@serverX vacation]$ ls -l
```

30. Chapters 5 and 6 may need a plot change. To prevent these changes from modifying original files, copy both files into **plot_change**. Move up one directory to **vacation**'s parent directory, then use one command from there.

```
31. [student@serverX vacation]$ cd ..
32. [student@serverX my_bestseller]$ cp
    chapters/mystery_chapter[56].odf plot_change
33. [student@serverX my_bestseller]$ ls -l chapters
```

```
[student@serverX my_bestseller]$ ls -l plot_change
```

34. To track changes, make three backups of chapter 5. Change to the **plot_change** directory. Copy **mystery_chapter5.odf** as a new file name to include the full date (Year-Mo-Da). Make another copy appending the current timestamp (as the number of seconds since the *epoch*) to ensure a unique file name. Also make a copy appending the current user (*\$USER*) to the file name. See the solution for the syntax of any you are unsure of (like what arguments to pass the **date**).

```
35. [student@serverX my_bestseller]$ cd plot_change
36. [student@serverX plot_change]$ cp
    mystery_chapter5.odf mystery_chapter5_$(date
    +%F).odf
37. [student@serverX plot_change]$ cp
    mystery_chapter5.odf mystery_chapter5_$(date
    +%s).odf
38. [student@serverX plot_change]$ cp
    mystery_chapter5.odf mystery_chapter5_$(USER).odf
```

```
[student@serverX plot_change]$ ls -l
```

Note, we could also make the same backups of the chapter 6 files too.

39. The plot changes were not successful. Delete the **plot_change** directory. First, delete all of the files in the **plot_change** directory. Change directory up one level because the directory cannot be deleted while it is the working directory. Try to delete the directory using the **rm** command *without* the *recursive* option. This attempt should fail. Now use the **rmdir** command, which will succeed.

```
40. [student@serverX plot_change]$ rm mystery*
41. [student@serverX plot_change]$ cd ..
42. [student@serverX my_bestseller]$ rm plot_change
43. rm: cannot remove 'plot_change': Is a directory
44. [student@serverX my_bestseller]$ rmdir plot_change
```

```
[student@serverX my_bestseller]$ ls -l
```

45. When the vacation is over, the **vacation** directory is no longer needed. Delete it using the **rm** command with the *recursive* option.

When finished, return to the home directory.

```
[student@serverX my_bestseller]$ rm -r vacation
[student@serverX my_bestseller]$ ls -l
[student@serverX my_bestseller]$ cd
```

Chapter 3. Getting Help in Red Hat Enterprise Linux

- 3.1. Reading Documentation Using man Command
- 3.2. Practice: Using the man Command
- 3.3. Reading Documentation Using pinfo Command
- 3.4. Practice: Using the pinfo Command
- 3.5. Reading Documentation in /usr/share/doc
- 3.6. Practice: Viewing Package Documentation
- 3.7. Getting Help From Red Hat
- 3.8. Practice: Creating and Viewing an SoS Report
- 3.9. Lab: Viewing and Printing Help Documentation

Introduction

Overview	
Goal	To resolve problems by using on-line help systems and Red Hat support utilities.
Objectives	<ul style="list-style-type: none">o Use the man Linux manual reader.o Use the pinfo GNU Info reader.o Use the Red Hat Package Manager (RPM) package documentation.o Use the redhat-support-tool command.
Sections	<ul style="list-style-type: none">o Reading Documentation Using man Command (and Practice)o Reading Documentation Using pinfo Command (and Practice)o Reading Documentation in /usr/share/doc (and Practice)o Getting Help From Red Hat (and Practice)
Lab	<ul style="list-style-type: none">o Viewing and Printing Help Documentation

3.1. Reading Documentation Using man Command

Objectives

After completing this section, students should be able to locate documentation and research answers about commands.

Introducing the man command

Reading documentation using man command

The historical Linux Programmer's Manual, from which man pages originate, was large enough to be multiple printed books. Each contained information for specific types of files, which have become the sections listed below. Articles are referred to as topics, as pages no longer applies.

□

Sections of the Linux manual

Section	Content type
1	User commands (<i>both executable and shell programs</i>) 用户可以操作的命令或可执行文件
2	System calls (<i>kernel routines invoked from user space</i>) 系统内核可调用的函数与工具等
3	Library functions (<i>provided by program libraries</i>) 常用的函数与函数库
4	Special files (<i>such as device files</i>) 设备文件的说明
5	File formats (<i>for many configuration files and structures</i>) 配置文件和其格式
6	Games (<i>historical section for amusing programs</i>) 娱乐、游戏
7	Conventions, standards, and miscellaneous (<i>protocols, file systems</i>) 约定与标准, 如Linux标准文件系统、网络协议
8	System administration and privileged commands (<i>maintenance tasks</i>) 系统管理员可用的管理命令
9	Linux kernel API (<i>internal kernel calls</i>) 内核的API

Note

Manual section 9 is a recent addition to Linux. Not all man section listings reference it.

To distinguish identical topic names in different sections, man page references include the section number in parentheses after the topic. For example, **passwd(1)** describes the command to change passwords, while **passwd(5)** explains the **/etc/passwd** file format for storing local user accounts.

To read specific man pages, use **man topic**. Topic contents display one screen at a time. Use **arrow** keys for **single line** scrolling or the **space bar** for the next screen. The **man** command searches manual sections in a configured order, **displaying popular sections first**. For example, **man passwd** displays **passwd(1)** by default. **To display the man page topic from a specific section**, include the section number argument: **man 5 passwd** displays **passwd(5)**.

Identify man pages by keyword

The ability to efficiently search for topics and navigate man pages is a critical administration skill. The following table lists basic man navigation commands:

□

Navigating man pages

Command	Result
---------	--------

Command	Result
Spacebar	Scroll forward (down) one screen
PageDown	Scroll forward (down) one screen
PageUp	Scroll backward (up) one screen
DownArrow	Scroll forward (down) one line
UpArrow	Scroll back (up) one line
d	Scroll forward (down) one half-screen
u	Scroll backward (up) one half-screen
/string	Search forward (down) for <i>string</i> in the man page
n	Repeat previous search forward (down) in the man page
N	Repeat previous search backward (up) in the man page
g	Go to start of the man page.
G	Go to end of the man page.
q	Exit man and return to the command shell prompt

Important

When performing searches, *string* allows **regular expression syntax**. While simple text (such as **passwd**) works as expected, **regular expressions use meta-characters** (such as **\$**, *****, **.**, and **^**) for more sophisticated pattern matching. Therefore, searching with strings which include program expression meta-characters, such as **make \$\$\$**, might yield unexpected results. Regular expressions and syntax are **discussed** in *Red Hat System Administration II*, and in the **regex(7)** man topic. [讨论](#)

▮ Searching for man pages by keyword

A keyword search of man pages is performed using **man -k keyword**, which displays a list of keyword-matching man page topics with section numbers.

```
[student@desktopX ~]$ man -k passwd
checkPasswdAccess (3) - query the SELinux policy database in the kernel.
chpasswd (8)          - update passwords in batch mode
ckpasswd (8)         - nnrpd password authenticator
fgetpwent_r (3)      - get passwd file entry reentrantly
getpwent_r (3)      - get passwd file entry reentrantly
...
passwd (1)           - update user's authentication tokens
sslpwssd (1ssl)     - compute password hashes
passwd (5)           - password file
passwd.nntp (5)     - Passwords for connecting to remote NNTP servers
```

```
passwd2des (3)      - RFS password encryption
...
```

Popular system administration topics are in [sections 1 \(user commands\)](#), [5 \(file formats\)](#), and [8 \(administrative commands\)](#). Administrators using certain troubleshooting tools also use [section 2](#) (system calls). The remaining sections are commonly for programmer reference or advanced administration.

Note

Keyword searches rely on an index generated by the **mandb(8)** command, which must be run as root. The command runs daily through **cron.daily**, or by **anacrontab** within an hour of boot if out of date.

Important

The **man** command **-K** option performs a full-text page search, not just titles and descriptions like the **-k**. A full-text search can use greater systems resources and take more time.

References

man(1), **mandb(8)**, **man-pages(7)**, **less(1)**, **intro(1)**, **intro(2)**, **intro(5)**, **intro(7)**, **intro(8)** man pages

3.2. Practice: Using the man Command

Guided exercise

In this lab, you will practice finding relevant information by using **man** options and arguments.

Outcomes

Familiarity with the **man** Linux manual system and practice finding useful information by searching and browsing.

Before you begin...

Perform the following steps on serverX unless directed otherwise.

1. View the **gedit(1)** man page.

```
[student@serverX ~]$ man 1 gedit
```

2. Research how to edit a specific file using **gedit** from the command line.

```
gedit filename
```

3. Research the **gedit** option used to begin an editing session with the cursor at the end of the file.

```
gedit + filename
```

4. Research the **su(1)** man page.

```
[student@serverX ~]$ man 1 su
```

5. Research what **su** does when the *username* argument is omitted.

```
su assumes a username of root.
```

6. Research how **su** behaves when a single dash option is used.

su starts a child *login shell* (creating login environment by sourcing login scripts). Without the single dash, a non-login child shell is created, matching the user's current environment.

7. Consult the **passwd(1)** man page. Determine the options that will lock and unlock a user account when this command is used by **root**.

```
[student@serverX ~]$ man 1 passwd
```

```
passwd -l username
```

```
passwd -u username
```

8. Locate the two principles to remember according to the **passwd man** page authors. Search for the word "principle".
 - Protect your password.
 - Choose a hard-to-guess password.
9. Consult the man page documenting the syntax of the **/etc/passwd** file. What is stored in the third field of each line?

The relevant man page is **passwd(5)**, found with **man -f passwd**.

The UID (numeric user ID) for each account.

10. Which command will list detailed information about a **zip** archive?

```
zipinfo(1) found with man -k zip
```

11. Which man page contains a list of parameters that can be passed to the kernel at boot time?

```
bootparam(7) found with man -k boot
```

12. Which command is used to tune **ext4** file system parameters?

```
tune2fs(8) found with man -k ext4
```

3.3. Reading Documentation Using pinfo Command

Objectives

After completing this section, students should be able to research answers using GNU Info documentation.

Introducing GNU info

Reading documentation using pinfo command

正式的

Man pages have a **formal** format useful as a command reference, but less useful as general documentation. For such documents, the GNU Project developed a different online documentation system, known as **GNU info**. Info documents are an important resource on a Red Hat Enterprise Linux system because many fundamental components and utilities, such as the *coreutils* package and *glibc* standard libraries, are either developed by the GNU Project or utilize the info document system.

基本的工具

```

File: dir      Node: Top      This is the top of the INFO tree

This (the Directory node) gives a menu of major topics.
Typing "q" exits, "?" lists all Info commands, "d" returns here,
"h" gives a primer for first-timers,
"mEmacs<Return>" visits the Emacs topic, etc.

In Emacs, you can click mouse button 2 on a menu item or cross reference
to select it.

* Menu:

Archiving
* Cpio: (cpio).          Copy-in-copy-out archiver to tape or disk.
* Tar: (tar).           Making tape (or disk) archives.

Basics
* Common options: (coreutils)Common options.
                        Common options.
* Coreutils: (coreutils).  Core GNU (file, text, shell) utilities.
* Date input formats: (coreutils)Date input formats.
* File permissions: (coreutils)File permissions.
                        Access modes.
* Finding files: (find).   Operating on files matching certain criteria.
* ed: (ed).               The GNU Line Editor.
Viewing line 25/2002, 1%

```

Figure 3.1: pinfo Info document viewer, top directory

Info documentation is structured as **hyperlinked info nodes**. This format is more **flexible** than man pages, allowing thorough discussion of complex commands and concepts. Like man pages, info nodes are read from the command line, using either the **info** or **pinfo** commands. Some commands and utilities have both **man** pages and info documentation; usually, the info documentation will be **more in-depth**. Compare the differences in **tar** documentation using **man** and **pinfo**: 更深入

```

[student@desktopX ~]$ man tar
[student@desktopX ~]$ pinfo tar

```

The **pinfo** info reader is more advanced than the original **info** command. Designed to match the **lynx** text web browser keystrokes, it also adds color. Info nodes for a particular topic are browsed with **pinfo topic**. Enter only **pinfo** for the info topics directory. New documentation nodes become available in **pinfo** when their corresponding software packages are installed.

GNU Info vs. man page navigation

The **info** command uses different navigation keystrokes than does **man**. The **info** command was designed to match the keystrokes of the hypertext-aware **lynx** web browser. Compare the key bindings in the following table:

□

pinfo and **man**, key binding comparison

Navigation	pinfo	man
------------	-------	-----

Navigation	pinfo	man
Scroll forward (down) one screen	PageDown or Space	PageDown or Space
Scroll backward (up) one screen	PageUp or b	PageUp or b
Display the directory of topics	d	–
Scroll forward (down) one half-screen	–	d
Display the parent node of a topic	u	–
Display the top (up) of a topic	HOME	1G
Scroll backward (up) one half-screen	–	u
Scroll forward (down) to next hyperlink	DownArrow	–
Open topic at cursor location	Enter	–
Scroll forward (down) one line	–	DownArrow or Enter
Scroll backward (up) to previous hyperlink	UpArrow	–
Scroll backward (up) one line	–	UpArrow
Search for a pattern	<i>/string</i>	<i>/string</i>
Display next node (chapter) in topic	n	–
Repeat previous search forward (down)	/ then Enter	n
Display previous node (chapter) in topic	p	–
Repeat previous search backward (up)	–	N
Quit the program	q	q

References

pinfo info (*Info: An Introduction*)

- o All Sections

pinfo pinfo (*Documentation for 'pinfo'*)

- o All Sections

The GNU Project <http://www.gnu.org/gnu/thegnuproject.html>

pinfo(1) and **info**(1) man pages

3.4. Practice: Using the pinfo Command

Guided exercise

In this lab, you will browse GNU Info documentation using command-line tools.

Outcomes

Understand program documentation in the GNU Info node system.

Before you begin...

Perform the following steps on serverX unless directed otherwise.

1. Invoke **pinfo** without arguments.

```
[student@serverX ~]$ pinfo
```

2. Navigate to the **Common options** topic.

Use **UpArrow** or **DownArrow** until **(coreutils) Common options** is highlighted. Press **Enter** to view this topic.

3. Browse through this **info** topic. Learn if long-style options can be abbreviated.

Use **PageUp** and **PageDown** to navigate through the topic. Yes, many programs allow long options to be abbreviated.

4. Determine what the symbols **--** signify when used as a command argument.

The symbols **--** signify the end of command *options* and the start of command *arguments* in complex commands where the shell's command-line parser might not correctly make the distinction.

5. Without exiting **pinfo**, move up to the **GNU Coreutils** node.

Press **u** to move up to the top node of the topic.

6. Move up again, to the top topic.

Press **u** again. Observe that when positioned at the top of a topic node, moving up returns to the directory of topics. Alternately, pressing **d** from any level or topic will move directly to the directory of topics.

7. Search for the pattern **nano** and select that topic.

Press **/** followed by the search pattern "nano". With the topic highlighted, press **Enter**.

8. In the **Introduction**, locate and select **Command Line Options**. Browse the topic.

Press **Enter** to select **Introduction**, then **DownArrow** and **Enter** to select **Command Line Options**. Use the arrow keys to browse the topic.

9. Move up one level to return to **Introduction**. Move to the next topic.

Press **u** to move up one level. The new location will be **nano's topic 1 Introduction**. Now press **n**. You will have moved to **nano's topic 2 Editor Basics**.

10. Exit **pinfo**.

Press **q** to quit **pinfo**.

11. Invoke **pinfo** again, specifying **nano** as the destination topic from the command line.

```
[student@serverX ~]$ pinfo nano
```

12. Select the **Editor Basics** topic.

Press **DownArrow** to highlight **Editor Basics**, then press **Enter** to select this topic.

13. Read the **Entering Text** and **Special Functions** subtopics.

Use arrow keys to highlight a topic, **PageUp** and **PageDown** to browse the text, then press **u** to move up one level. Press **q** to quit **pinfo** when you are finished.

3.5. Reading Documentation in `/usr/share/doc`

Objectives

After completing this section, students should be able to research information using Red Hat Package Manager documentation.

Introducing package documentation

Reading documentation in `/usr/share/doc`

In addition to **man** and **pinfo**, developers may also choose to include documentation in their application's RPM distribution package. When the package is installed, files recognized as documentation are moved to `/usr/share/doc/packagename`. Software package builders may include anything deemed helpful as a complement to, but not duplicating, **man** pages. GNU packages also use `/usr/share/doc` to supplement info nodes.

Most packages include files describing package distribution licensing. Some packages include extensive PDF- or HTML-based documentation. Accordingly, a useful package browsing method is pointing a browser of choice to `file:///usr/share/doc` and utilizing a mouse.

```
[student@desktopX ~]$ firefox file:///usr/share/doc
```

Some packages come with extensive examples, configuration file templates, scripts, tutorials, or user guides. Browse `/usr/share/doc/vsftpd-*` as an example. Some documentation is sparse; the `zip` utility includes the compression algorithm used and little else. Other packages includes large user manuals or developer guides, or electronic copies of related, published books.

Note

Developers may choose to bundle extensive documentation in a separate RPM. The `gnuplot` program has the extra `gnuplot-doc` package, which must be installed separately. Other similar packages to browse include `bash-doc` and `samba-doc`. Often, extra packages are found in Red Hat Enterprise Linux's *Optional* software channel.

Many packages also include developer documentation, such as an Application Programming Interface (API) specification, provided in a package with a name ending in `-devel` or similar. Packages may include additional files, such as headers; useful documentation usually only needed for software development or compiling.

Note

The kernel itself has a significant documentation package. The *kernel-doc* package is a treasure of kernel, driver, tuning, and advanced configuration information. Experienced system administrators regularly research *kernel-doc* files.

References

hier(7) man page

- Discusses the hierarchy of Linux directories, including `/usr/share/doc`.

3.6. Practice: Viewing Package Documentation

Guided exercise

In this lab, you will research the documentation under `/usr/share/doc` to answer questions. Use your choice of **less**, **gedit**, or a browser to view the documentation file contents.

Outcomes

More familiarity, through practice, with the types of information that developers include with their software packages.

Before you begin...

Perform the following steps on serverX unless directed otherwise.

1. Where can you find the latest news about the *vim* project?

```
2. [student@serverX ~]$ cd /usr/share/doc
```

```
[student@serverX doc]$ less vim-common-*/README.txt
```

View the *vim-common* README and search for “news”.

3. What is the wiki URI for the *yum* package?

```
[student@serverX doc]$ less yum-3*/README
```

Search for “wiki” in `/usr/share/doc/yum-3*/README`.

4. What examples are provided for the command-line **bc** calculator?

```
[student@serverX doc]$ ls -l bc-*/Examples
```

Found in the directory `/usr/share/doc/bc-*/Examples`.

5. How would you read the provided GRUB2 manual?

```
[student@serverX doc]$ firefox grub2-tools-*/grub.html
```

Use **firefox** to

display `/usr/share/doc/grub2-tools-*/grub.html`.

6. What software provides its document as a separate package?

```
7. [student@serverX doc]$ yum list *-doc*
```

```
8. [student@serverX doc]$ cd
```

```
[student@serverX ~]$
```

Use **yum** to display only those packages that contain “-doc”, “-docs”, or “-documentation” in the package name. When finished, return to the home directory.

3.7. Getting Help From Red Hat

Objectives

After completing this section, students should be able to view Knowledgebase information and manage support cases from the command line.

Red Hat Customer Portal

Red Hat Customer Portal (<https://access.redhat.com>) provides customers with access to everything provided with their subscription through one convenient location. Customers can search for solutions, FAQs, and articles through Knowledgebase. Access to official product documentation is provided. Support tickets can be submitted and managed. Subscriptions to Red Hat products can be attached to and detached from registered systems, and software downloads, updates, and evaluations can be obtained. Parts of the site are accessible to everyone, while others are exclusive to customers with active subscriptions. Help with getting access to Customer Portal is available at <https://access.redhat.com/help/>.

Customers can work with Red Hat Customer Portal through a web browser. This section will introduce a command line tool that can also be used to access Red Hat Customer Portal services, **redhat-support-tool**.

Knowledgebase




Figure 3.2: Knowledgebase at the Red Hat Customer Portal

Using redhat-support-tool to search Knowledgebase

Getting help from Red Hat

The Red Hat Support Tool utility **redhat-support-tool** provides a text console interface to the subscription-based Red Hat Access services. Internet access is required to reach the Red Hat Customer Portal. The **redhat-support-tool** is text-based for use from any terminal or SSH connection; no graphical interface is provided.

The **redhat-support-tool** command may be used as an interactive shell or invoked as individually executed commands with options and arguments. The tool's available syntax is identical for both methods. By default, the program launches in shell mode. Use the provided **help** sub-command to see all available commands. Shell mode supports tab completion and the ability to call programs in the parent shell.

```
[student@desktopX ~]$ redhat-support-tool
Welcome to the Red Hat Support Tool.
Command (? for help):
```

When first invoked, **redhat-support-tool** prompts for required Red Hat Access subscriber login information. To avoid repetitively supplying this information, the tool asks to store account information in the user's home directory

(**\$HOME/.redhat-support-tool/redhat-support-tool.conf**). If a Red Hat Access account is shared by many users, the **--global** option can save account information to **/etc/redhat-support-tool.conf**, along with other systemwide configuration.

The tool's **config** command modifies tool configuration settings.

The **redhat-support-tool** allows subscribers to search and display the same Knowledgebase content seen when on the Red Hat Customer Portal. Knowledgebase permits keyword searches, similar to the **man** command. Users can enter error codes, syntax from log files, or any mix of keywords to produce a list of relevant solution documents.

The following is an initial configuration and basic search demonstration:

```
[student@desktopX ~]$ redhat-support-tool
Welcome to the Red Hat Support Tool.

Command (? for help): search How to manage system entitlements with
subscription-manager

Please enter your RHN user ID: subscriber
Save the user ID in /home/student/.redhat-support-tool/redhat-support-tool.conf
(y/n): y

Please enter the password for subscriber: password
Save the password for subscriber in
/home/student/.redhat-support-tool/redhat-support-tool.conf (y/n): y
```

After prompting the user for the required user configuration, the tool continues with the original search request:

```
Type the number of the solution to view or 'e' to return to the previous menu.
 1 [ 253273:VER] How to register and subscribe a system to Red Hat Network
   (RHN) using Red Hat Subscription Manager (RHSM)?
 2 [ 17397:VER] What are Flex Guest Entitlements in Red Hat Network?
 3 [ 232863:VER] How to register machines and manage subscriptions using Red
   Hat Subscription Manager through an invisible HTTP proxy / Firewall?
3 of 43 solutions displayed. Type 'm' to see more, 'r' to start from the beginning
again, or '?' for help with the codes displayed in the above output.

Select a Solution:
```

Specific sections of solution documents may be selected for viewing.

```
Select a Solution: 1

Type the number of the section to view or 'e' to return to the previous menu.
 1 Title
```

```

2 Issue
3 Environment
4 Resolution
5 Display all sections
End of options.
Section: 1

Title
=====
===
How to register and subscribe a system to Red Hat Network (RHN) using Red Hat
Subscription Manager (RHSM)?
URL:      https://access.redhat.com/site/solutions/253273
(END) q
[student@desktopX ~]$

```

Directly access Knowledgebase articles by document ID

Locate online articles directly using the tool's **kb** command with the Knowledgebase document ID. Returned documents scroll on the screen without pagination, allowing a user to redirect the output using other local commands. This example views the document with the **less** command:

```

[student@desktopX ~]$ redhat-support-tool kb 253273 | less

Title:  How to register and subscribe a system to Red Hat Network (RHN) using
Red Hat Subscription Manager (RHSM)?
ID:     253273
State:  Verified: This solution has been verified to work by Red Hat Customers
and Support Engineers for the specified product version(s).
URL:    https://access.redhat.com/site/solutions/253273
: q

```

Documents retrieved in unpaginated format are easy to send to a printer, convert to PDF or other document format, or to redirect to a data entry program for an incident tracking or change management system, using other utilities installed and available in Red Hat Enterprise Linux.

Using redhat-support-tool to manage support cases

One benefit of a product subscription is access to technical support through Red Hat Customer Portal. Depending on the system's subscription support level, Red Hat may be contacted through on-line tools or by phone.

See https://access.redhat.com/site/support/policy/support_process for links to detailed information about the support process.

Preparing a bug report

Before contacting Red Hat Support, gather relevant information for a bug report.

Define the problem. Be able to clearly state the problem and its symptoms. Be as specific as possible.

Detail the steps which will reproduce the problem.

Gather background information. Which product and version is affected? Be ready to provide relevant diagnostic information. This can include output of **sosreport**, discussed later in this section. For kernel problems, this could include the system's **kdump** crash dump or a digital photo of the kernel backtrace displayed on the monitor of a crashed system.

Determine the severity level. Red Hat uses four severity levels to classify issues. *Urgent* and *High* severity problem reports should be followed by a phone call to the relevant local support center

(see <https://access.redhat.com/site/support/contact/technicalSupport>).

Severity	Description
<i>Urgent</i> (Severity 1)	A problem that severely impacts your use of the software in a production environment (such as loss of production data or in which your production systems are not functioning). The situation halts your business operations and no procedural workaround exists.
<i>High</i> (Severity 2)	A problem where the software is functioning but your use in a production environment is severely reduced. The situation is causing a high impact to portions of your business operations and no procedural workaround exists.
<i>Medium</i> (Severity 3)	A problem that involves partial, non-critical loss of use of the software in a production environment or development environment. For production environments, there is a medium-to-low impact on your business, but your business continues to function, including by using a procedural workaround. For development environments, where the situation is causing your project to no longer continue or migrate into production.
<i>Low</i> (Severity 4)	A general usage question, reporting of a documentation error, or recommendation for a future product enhancement or modification. For production environments, there is low-to-no impact on your business or the performance or functionality of your system. For development environments, there is a medium-to-low impact on your business, but your business continues to function, including by using a procedural workaround.

Managing a bug report with **redhat-support-tool**

Subscribers may create, view, modify, and close Red Hat Support cases

using **redhat-support-tool**. When support cases are opened or maintained, users may include files or documentation, such as diagnostic reports (sosreport). The tool uploads and attaches files to online cases. Case details including *product*, *version*, *summary*, *description*, *severity*, and *case group* may be assigned with command options or letting the tool prompt for required information. In the following example, the **--product** and **--version** options are specified,

but **redhat-support-tool** would provide a list of choices for those options if the **opencase** command did not specify them.

```
[student@desktopX ~]$ redhat-support-tool
Welcome to the Red Hat Support Tool.

Command (? for help): opencase --product="Red Hat Enterprise Linux"
--version="7.0"

Please enter a summary (or 'q' to exit): System fails to run without power
Please enter a description (Ctrl-D on an empty line when complete):
When the server is unplugged, the operating system fails to continue.

1  Low
2  Normal
3  High
4  Urgent

Please select a severity (or 'q' to exit): 4

Would you like to assign a case group to this case (y/N)? N

Would see if there is a solution to this problem before opening a support case?
(y/N) N

-----
---

Support case 01034421 has successfully been opened.
```

Including diagnostic information by attaching a SoS report archive

Including diagnostic information when a support case is first created contributes to quicker problem resolution. The **sosreport** command generates a compressed tar archive of diagnostic information gathered from the running system. The **redhat-support-tool** prompts to include one if an archive has been created previously:

```
Please attach a SoS report to support case 01034421. Create a SoS report as
the root user and execute the following command to attach the SoS report
directly to the case:
```

```
redhat-support-tool addattachment -c 01034421 path to sosreport
```

```
Would you like to attach a file to 01034421 at this time? (y/N) N
```

```
Command (? for help):
```

If a current SoS report is not already prepared, an administrator can generate and attach one later, using the tool's **saddattachment** command as advised previously. This section's practice exercise will provide the steps for creating and viewing a current SoS diagnostic report.

Support cases can also be viewed, modified, and closed by you as the subscriber:

```
Command (? for help): listcases
```

```
Type the number of the case to view or 'e' to return to the previous menu.
```

```
1 [Waiting on Red Hat] System fails to run without power
```

```
No more cases to display
```

```
Select a Case: 1
```

```
Type the number of the section to view or 'e' to return to the previous menu.
```

```
1 Case Details
```

```
2 Modify Case
```

```
3 Description
```

```
4 Recommendations
```

```
5 Get Attachment
```

```
6 Add Attachment
```

```
7 Add Comment
```

```
End of options.
```

```
Option: q
```

```
Select a Case: q
```

```
Command (? for help):q
```

```
[student@desktopX ~]$ redhat-support-tool modifycase  
--status=Closed 01034421
```

```
Successfully updated case 01034421
```

```
[student@desktopX ~]$
```

The Red Hat Support Tool has advanced application diagnostic and analytic capabilities. Using kernel crash dump core files, **redhat-support-tool** can create and extract a *backtrace*, a report of the active stack frames at the point of a crash dump, to provide onsite diagnostics and open a support case.

The tool also provides log file analysis. Using the tool's **analyze** command, log files of many types, including operating system, JBoss, Python, Tomcat, oVirt, and others, can be parsed to recognize problem symptoms, which can then be viewed and diagnosed individually. Providing preprocessed analysis, as opposed to raw data such as crash dump or log files, allows support cases to be opened and made available to engineers more quickly.

References

sosreport(1) man page

[Red Hat Access: Red Hat Support Tool](#)

[Red Hat Support Tool First Use](#)

3.8. Practice: Creating and Viewing an SoS Report

Guided exercise

In this lab, you will use the `sosreport` command to generate a SoS report, then view the contents of that diagnostic archive.

Outcomes

A compressed tar archive of systemwide diagnostic information.

Before you begin...

Perform the following steps on `serverX` unless directed otherwise.

1. If currently working as a non-root user, switch to root.

```
2. [student@serverX ~]$ su -
```

```
Password: redhat
```

3. Run the `sosreport` command. This may take many minutes on larger systems.

```
4. [root@serverX ~]# sosreport
```

```
5.
```

```
6. sosreport (version 3.0)
```

```
7.
```

```
8. This command will collect system configuration and
```

```
9. diagnostic information from this Red Hat Enterprise Linux
```

```
10. system. An archive containing the collected information
```

```
11. will be generated in /var/tmp and may be provided to a Red
```

```
12. Hat support representative or used for local diagnostic or
```

```
13. recording purposes.
```

```
14.
```

```
15. Any information provided to Red Hat will be treated in
```

```
16. strict confidence in accordance with the published support
```

```
17. policies at:
```

```
18.
```

```
19. https://access.redhat.com/support/
```

```
20.
```

```
21. The generated archive may contain data considered
```

```
22. sensitive and its content should be reviewed by the
```

```
23. originating organization before being passed to any third party.
```

```
24.
```


25. No changes will be made to system configuration.
- 26.
27. Press ENTER to continue, or CTRL-C to quit. **ENTER**
- 28.
29. Please enter your first initial and last name [serverX.example.com]: **yourname**

Please enter the case number that you are generating this report for: **01034421**

Press **Enter**. Provide the requested information. Make up a value for the case number.

```
Running 17/74: general...
Creating compressed archive...

Your sosreport has been generated and saved in:

  /var/tmp/sosreport-yourname.01034421-20140129000049.tar.xz

The checksum is: b2e78125290a4c791162e68da8534887

Please send this file to your support representative.
```

30. Change directory to **/var/tmp**, and unpack the archive.

```
31. [root@serverX ~]# cd /var/tmp
```

```
[root@serverX tmp]# tar -xvJf sosreport-*.tar.xz
```

32. Change directory to the resulting subdirectory and browse the files found there.

```
33. [root@serverX ~]# cd
    sosreport-yourname.01034421-20140129000049
```

```
[root@serverX
sosreport-yourname.01034421-20140129000049]# ls -lR
```

Open files, list directories, and continue to browse to become familiar with the information included in SoS reports. In the form of the original archived and compressed file, this is the diagnostic information you would be attaching to a **redhat-support-tool** support case. When finished, remove the archive directory and files and return to your home directory.

```
[root@serverX
sosreport-yourname.01034421-20140129000049]# cd
/var/tmp
[root@serverX tmp]# rm -rf sosreport*
```

```
[root@serverX tmp]# exit
[student@serverX ~]$
```

3.9. Lab: Viewing and Printing Help Documentation

Performance checklist

In this lab, you will practice research methods typically used by system administrators to learn how to perform necessary tasks.

Outcomes

- Accomplish a given task; practice locating relevant commands by searching **man** pages and **pinfo** nodes.
- Learn new options for commonly used documentation commands.
- Recognize various document file formats; use appropriate tools to view and print documentation and other non-text formatted files.

Before you begin...

Perform the following steps on serverX unless directed otherwise.

1. Research **man(1)** to determine how to prepare a man page for printing. What format or rendering language is commonly used?

```
[student@serverX ~]$ man man
```

man uses **-t** to prepare a man page for printing, using PostScript.

2. Create a formatted output file of the **passwd** man page. Determine the file content format.

```
3. [student@serverX ~]$ man -t passwd > passwd.ps
```

```
4. [student@serverX ~]$ file passwd.ps
```

```
[student@serverX ~]$ less passwd.ps
```

The file is in PostScript format, learned using the **file** command and confirmed by viewing the file contents. Notice the header lines of PostScript information.

5. Research using **man** to learn the command(s) used for viewing or printing PostScript files.

```
[student@serverX ~]$ man -k postscript viewer
```

Using multiple words with the **-k** option finds man pages matching *either* word; those with "postscript" *or* "viewer" in their descriptions. Notice the **evince(1)** and **ghostscript(1)** (or **gs(1)**) commands in the output.

6. Research **evince(1)** using **man** to learn how to use the viewer in preview mode. Also, determine how to open a document starting on a specific page.

```
[student@serverX ~]$ man evince
```

The **-w** (or **--preview**) option opens **evince** in preview mode. The **-i** option is used to specify a starting page.

7. View your PostScript file using the various **evince** options you researched. Close your document file when you are finished.

```
8. [student@serverX ~]$ evince passwd.ps
9. [student@serverX ~]$ evince -w passwd.ps
```

```
[student@serverX ~]$ evince -i 3 passwd.ps
```

While normal **evince** mode allows full-screen and presentation-style viewing, the **evince** preview mode is useful for quick browsing and printing. Notice the **print icon** at the top.

10. Using **man**, research **lp(1)** to determine how to print any document starting on a specific page. Without actually entering any commands (since there are no printers), what would be the syntax, on one command line, to print only pages 2 and 3 of your PostScript file?

```
[student@serverX ~]$ man lp
```

One answer is **lp passwd.ps -P 2-3**.

From **lp(1)**, learn that the **-P** option specifies pages. The **lp** command spools to the *default* printer, sending only the page range starting on 2 and ending on 3.

Note

There are currently no printers configured in the classroom. However, you may practice later using printer models configured in your own environment. Familiarity with these commands is often useful.

11. Using **pinfo**, look for GNU info about the **evince** viewer.

```
[student@serverX ~]$ pinfo evince
```

Notice that the **evince(1)** man page displays instead. The **pinfo** document viewer looks for the relevant man page when no appropriate GNU documentation node exists for the requested topic. Press **q** to close **pinfo**.

12. As an opportunity to observe the abundance of GNU fundamental utilities, use **pinfo** to locate and browse all document nodes for the *coreutils* commands and programs.

```
[student@serverX ~]$ pinfo
```

From the directory node, press **DownArrow** until the link is selected for **Coreutils: Core GNU (file, text, shell) utilities**. Press **Enter** to follow the link to **GNU Coreutils**. Notice the long menu listing, with **Introduction** currently selected. Press **Enter**. At the top of the screen, pay attention to the header, which displays the previous, current, and next nodes. Browse the information, press **n** for the next node, and repeat. Browse each screen, simply noticing the commands and their descriptions. Continue until node **29 Opening the Software Toolbox** is reached. Read this chapter in its entirety using the navigation

you have learned. When finished, return the way you came by using *only* **LeftArrow** until the top directory node is finally reached. Press **q** to close **pinfo**.

13. Using **firefox**, open the system's package documentation directory and browse into the **man-db** package subdirectory. View the provided manual(s).

```
[student@serverX ~]$ firefox /usr/share/doc
```

Remember that bookmarks can be made for any directories that are frequently used. After

Overview

browsing to the **man-db** directory, click to open and view the text version of the manual, then close it. Click to open the PostScript version. As observed earlier, **evince** is the system's default viewer for PostScript and PDF documents. You may wish to return to these documents later to become more knowledgeable about **man**. When finished, close the **evince** viewer.

14. Using the open **Firefox** browser, locate and browse into the **initscripts** package subdirectory. View the **sysconfig.txt** file, which describes important system configuration options stored in the **/etc/sysconfig** directory.

Notice how useful a browser is for locating and viewing local system documentation. Close the document when finished, but leave Firefox open.

□ Chapter 4. Creating, Viewing, and Editing Text Files

4.1. Redirecting Output to a File or Program

4.2. Practice: I/O Redirection and Pipelines

4.3. Editing Text Files from the Shell Prompt

4.4. Practice: Editing Files with Vim

4.5. Editing Text Files with a Graphical Editor

4.6. Practice: Copying Text Between Windows

4.7. Lab: Creating, Viewing, and Editing Text Files

Introduction

Goal	To create, view, and edit text files from command output or in an editor.
Objectives	<ul style="list-style-type: none"> ○ Redirect the text output of a program to a file or to another program. ○ Edit existing text files and create new files from the shell prompt with a text editor. ○ Copy text from a graphical window to a text file using a text editor running in the graphical environment.
Sections	<ul style="list-style-type: none"> ○ Redirecting Output to a File or Program (and Practice) ○ Editing Text Files from the Shell Prompt (and Practice) ○ Editing Text Files with a Graphical Editor (and Practice)
Lab	<ul style="list-style-type: none"> ○ Creating, Viewing, and Editing Text Files

4.1. Redirecting Output to a File or Program

Objectives

After completing this section, students should be able to:

- Describe the technical terms standard input, standard output, and standard error.
- Use redirection characters to control output to files.
- Use piping to control output to other programs.

Standard input, standard output, and standard error

Redirecting output to a file or program

A process structure is constructed with numbered channels (*file descriptors*) to manage open files.

Processes connect to files to reach data content or devices these files represent. Processes are created with default connections for channels 0, 1, and 2, known as *standard input*, *standard output*, and *standard error*. Processes use channels 3 and above to connect to other files.

Figure 4.1: Process I/O channels (file descriptors)

□

Channels (File Descriptors)

Number	Channel name	Description	Default connection	Usage
0	stdin	Standard input	Keyboard	read only
1	stdout	Standard output	Terminal	write only
2	stderr	Standard error	Terminal	write only
3+	filename	Other files	<i>none</i>	read and/or write

□ Redirecting output to a file

Channel *redirection* replaces default channel destinations with file names representing either output files or devices. Using redirection, process output and error messages can be captured as file contents, sent to a device, or discarded.

Redirecting **stdout** suppresses process output from appearing on the terminal. As seen in the following table, redirecting *only* **stdout** does not suppress **stderr** error messages from displaying on the terminal. The special file **/dev/null** quietly discards channel output redirected to it.

□

Output Redirection Operators

Usage	Explanation ^(note)	Visual aid
<code>>file</code>	redirect stdout to a file ⁽¹⁾	
<code>>>file</code>	redirect stdout to a file, append to current file content ⁽²⁾	
<code>2>file</code>	redirect stderr to a file ⁽¹⁾	
<code>2>/dev/null</code>	discard stderr error messages by redirecting to /dev/null	
<code>&>file</code>	combine stdout and stderr to one file ⁽¹⁾	
<code>>>file 2>&1</code>	combine stdout and stderr , append to current file content ⁽²⁾ ⁽³⁾	
Note:	<p>⁽¹⁾ Overwrite existing file, create file if new.</p> <p>⁽²⁾ Append existing file, create file if new.</p> <p>⁽³⁾ The order of redirection is important to avoid unexpected command behavior. 2>&1 sends stderr to the same place as stdout. For this to work, stdout needs to be redirected first, <i>before</i> adding stderr to stdout. Although &>> is an</p>	

Usage	Explanation ^(note)	Visual aid
	<p>alternate way to append both stdout and stderr to a file, 2>&1 is the method needed to send both stdout and stderr through a pipe.</p>	

Examples for output redirection

Many routine administration tasks are simplified by using redirection. Use the previous table to assist while considering the following examples:

- Save a timestamp for later reference.

```
[student@desktopX ~]$ date > /tmp/saved-timestamp
```

- Copy the last 100 lines from a log file to another file.

```
[student@desktopX ~]$ tail -n 100 /var/log/dmesg > /tmp/last-100-boot-messages
```

- Concatenate four files into one.

```
[student@desktopX ~]$ cat file1 file2 file3 file4 > /tmp/all-four-in-one
```

- List the home directory's hidden and regular file names into a file.

```
[student@desktopX ~]$ ls -a > /tmp/my-file-names
```

- Append output to an existing file.

- ```
[student@desktopX ~]$ echo "new line of information" >> /tmp/many-lines-of-information
```

```
[student@desktopX ~]$ diff previous-file current-file >> /tmp/tracking-changes-made
```

- In the next examples, errors are generated since normal users are denied access to system directories. Redirect errors to a file while viewing normal command output on the terminal.

```
[student@desktopX ~]$ find /etc -name passwd 2> /tmp/errors
```

- Save process output and error messages to separate files.

```
[student@desktopX ~]$ find /etc -name passwd > /tmp/output 2> /tmp/errors
```



- Ignore and discard error messages.

```
[student@desktopX ~]$ find /etc -name passwd > /tmp/output 2> /dev/null
```

- Store output and generated errors together.

```
[student@desktopX ~]$ find /etc -name passwd &> /tmp/save-both
```

- Append output and generated errors to an existing file.

```
[student@desktopX ~]$ find /etc -name passwd >> /tmp/save-both 2>&1
```

## ❏ Constructing pipelines

*Redirection* controls channel output to or from *files* while *piping* sends channel output to another *process*.

Figure 4.8: Process I/O piping

Examples for process piping redirection

Paginate a command's long output.

```
[student@desktopX ~]$ ls -l /usr/bin | less
```

Count the number of lines in an output or listing.

```
[student@desktopX ~]$ ls | wc -l > /tmp/how-many-files
```

Grab the first lines, last lines, or selected lines of command output.

```
[student@desktopX ~]$ ls -t | head -n 10 >
/tmp/ten-last-changed-files
```

*Figure 4.9: Process I/O piping with tee*

#### Examples for using the **tee** command for piping

The **tee** command displays or redirects the intermediate result normally suppressed due to piping. In the first example, the **ls** listing is viewed on a terminal while simultaneously being stored in a file.

```
[student@desktopX ~]$ ls -l | tee /tmp/saved-output
```

Determine the terminal device for the current window. Send the results as mail and view the same results in this window.

```
[student@desktopX ~]$ tty
```

```
/dev/pts/0
```

```
[student@desktopX ~]$ ls -l | tee /dev/pts/0 | mail -s subject student@desktop1.example.com
```

#### References

**info bash** (*The GNU Bash Reference Manual*)

- Section 3.2.2: Pipelines

**info bash** (*The GNU Bash Reference Manual*)

- Section 3.6: Redirections

**info coreutils 'tee invocation'** (*The GNU coreutils Manual*)

- Section 17.1: Redirect output to multiple files or processes

**bash(1), cat(1), head(1), less(1), mail(1), tee(1), tty(1), wc(1)** man pages

## 4.2. Practice: I/O Redirection and Pipelines

### Quiz

Match the following items to their counterparts in the table.

| Result needed                                                                | Redirection syntax used |
|------------------------------------------------------------------------------|-------------------------|
| Display command output to terminal, ignore all errors.                       | 2>/dev/null             |
| Send command output to file; errors to different file.                       | >file 2>file2           |
| Send output and errors to the same new, empty file.                          | &>file                  |
| Send output and errors to the same file, but preserve existing file content. | >>file 2>&1             |
| Run a command, but throw away all possible terminal displays.                | &>/dev/null             |
| Send command output to both the screen and a file at the same time.          | tee file                |
| Run command, save output in a file, discard error messages.                  | > file 2> /dev/null     |

## 4.3. Editing Text Files from the Shell Prompt

### Objectives

After completing this section, students should be able to:

- Create new files and edit existing text files from the shell prompt.
- Navigate within an editor to effectively accomplish editing tasks.

### Editing files with Vim

*Editing text files from the shell prompt*

A key design principle of Linux is that information is stored in text-based files. Text files include both *flat files* with rows of similar information, such as configuration files in **/etc**, and *Extensible Markup Language (XML)* files, which define data structure through text tags, seen in application configuration files throughout both **/etc** and **/usr**. The advantage of text files is that they can be

moved or shared between systems without requiring conversion, and can be viewed and edited using any simple text editor.

Vim is an improved version of the vi editor distributed with Linux and UNIX systems. Vim is highly configurable and efficient for practiced users, including such features as split screen editing, color formatting, and highlighting for editing text.

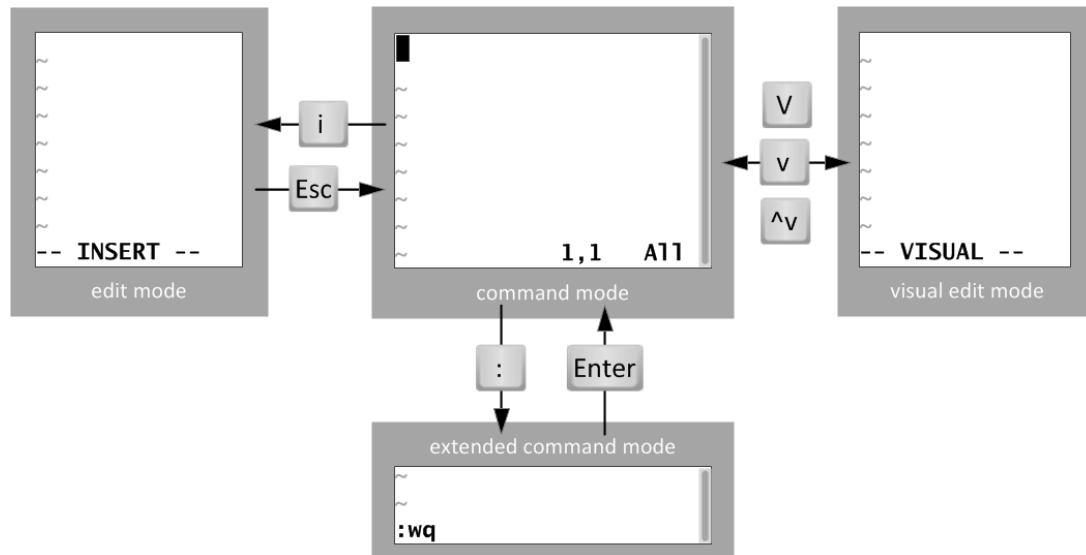


Figure 4.10: Moving between Vim modes

When first opened, Vim starts in *command mode*, used for navigation, cut and paste, and other text manipulation. Enter each of the other modes with single character keystrokes to access specific editing functionality:

- An **i** keystroke enters *insert mode*, where all text typed becomes file content. Pressing **Esc** returns to command mode.
- A **v** keystroke enters *visual mode*, where multiple characters may be selected for text manipulation. Use **V** for multi-line and **Ctrl-v** for block selection. The same keystroke used to enter visual mode (**v**, **V** or **Ctrl-v**) is used to exit.
- The **:** keystroke begins *extended command mode* for tasks like writing the file to save it, and quitting the Vim editor.

#### The minimum, basic Vim workflow

**Vim** has efficient, coordinated keystrokes for advanced editing tasks. Although considered useful with practice, Vim's capabilities can overwhelm new users. The following workflow presents the *minimum* keystrokes every Vim user must learn to accomplish *any* editing task.

The instructor will demonstrate a typical file editing session using only basic Vim keystrokes.

1. Open a file with **vim filename**.
2. Repeat this text entry cycle, as many times as the task requires:
  - Use arrow keys to position the cursor.
  - Press **i** to enter insert mode.
  - Enter text.
  - Press **Esc** to return to command mode.
  - If necessary, press **u** to undo mistaken edits on the current line.

3. Repeat this text deletion cycle, as many times as the task requires:
  - Use arrow keys to position the cursor.
  - Press **x** to delete a selection of text.
  - If necessary, use **u** to undo mistaken edits on the current line.
4. To save or exit, choose one of the following to write or discard file edits:
  - Enter **:w** to write (save) the file and remain in command mode for more editing.
  - Enter **:wq** to write the file and quit Vim.
  - Enter **:q!** to quit Vim, but discard all file changes since the last write.

#### Rearranging existing text

In **Vim**, copy and paste is known as *yank and put*, using command characters **y** and **p**. Begin by positioning the cursor on the first character to be selected, then enter visual mode. Use arrow keys to expand the visual selection. When ready, press **y** to *yank* the selection into memory. Position the cursor at the new location, then press **p** to put the selection at the cursor.

The instructor will demonstrate “yank and put” using visual mode.

1. Open a file with **vim filename**.
2. Repeat this text selection cycle, as many times as the task requires:
  - Use arrow keys to position the cursor to the first character.
  - Press **v** to enter visual mode.
  - Use arrow keys to position the cursor to the last character.
  - Press **y** to yank (copy) the selection.
  - Use arrow keys to position the cursor at the insert location.
  - Press **p** to put (paste) the selection.
3. To save or exit, choose one of the following to write or discard file edits:
  - Enter **:w** to write (save) the file and remain in command mode for more editing.
  - Enter **:wq** to write the file and quit Vim.
  - Enter **:q!** to quit Vim, but discard all file changes since the last write.

#### Note

Beware the advanced Vim user offering shortcuts and tricks before the basics are mastered. Vim requires practice to become efficient. It is recommended to continue learning new keystrokes to extend Vim's usefulness. For those curious how extensive this can be, perform an Internet search for “Vim tips”.

An in-depth presentation of **Vim** is included in the *Red Hat Enterprise Linux 7 System Administration II* course.

#### References

**vim**(1) man page

[Vim the editor](#)

## 4.4. Practice: Editing Files with Vim

### Guided exercise

In this lab, you will use a locally installed resource to practice entry-level **vim** editor techniques.

#### Outcomes:

Experience with vim, and knowledge about using **vimtutor** to gain competency.

#### Before you begin...

In this exercise, use the existing Vim tutorial bundled with the Vim editor. The installed *vim-enhanced* package provides **vimtutor**. For each exercise step, use the corresponding lesson in vimtutor to practice. Return here when the lesson step is complete. Perform the following steps on serverX unless directed otherwise.

1. Open **vimtutor**. Read the Welcome screen and perform *Lesson 1.1*.

```
[student@serverX ~]$ vimtutor
```

In the lecture, only keyboard arrow keys were used for navigation. In **vi**'s early years, users could not rely on working keyboard mappings for arrow keys. Therefore, **vi** was designed with commands using only standard character keys, such as the conveniently grouped **h**, **j**, **k**, and **l**. Here is one way to remember them:

**h**ang back, **j**ump down, **k**ick up, **l**eap forward.

2. Return to the **vimtutor** window. Perform *Lesson 1.2*.

This early lesson teaches how to quit without having to keep an unwanted file change. All changes are lost, but this is better than leaving a critical file in an incorrect state.

3. Return to the **vimtutor** window. Perform *Lesson 1.3*.

Vim has faster, more efficient keystrokes to delete an exact amount of words, lines, sentences, and paragraphs. However, any editing job *can* be accomplished using only **x** for single-character deletion.

4. Return to the **vimtutor** window. Perform *Lesson 1.4*.

The minimum required keystrokes are for entering and leaving edit mode, arrow keys, and deleting. For most edit tasks, the first key pressed is **i**.

5. (Optional) Return to the **vimtutor** window. Perform *Lesson 1.5*.

In the lecture, only the **i** (*insert*) command was taught as the keystroke to enter edit mode. This vimtutor lesson demonstrates that other keystrokes are available to change the cursor placement when insert mode is entered. However, once in insert mode, all text typed is still file content.

6. Return to the **vimtutor** window. Perform *Lesson 1.6*.

Save the file by **w**riting and **q**uitting. This is the last lesson for the minimum *required* keystrokes to be able to accomplish any editing task.

7. Return to the **vimtutor** window. Finish by reading the *Lesson 1 Summary*.

There are six more multi-step lessons in **vimtutor**. None are assigned as further lessons for this course, but feel free to use **vimtutor** on your own to learn more about **Vim**.

## 4.5. Editing Text Files with a Graphical Editor

### Objectives

After completing this section, students should be able to:

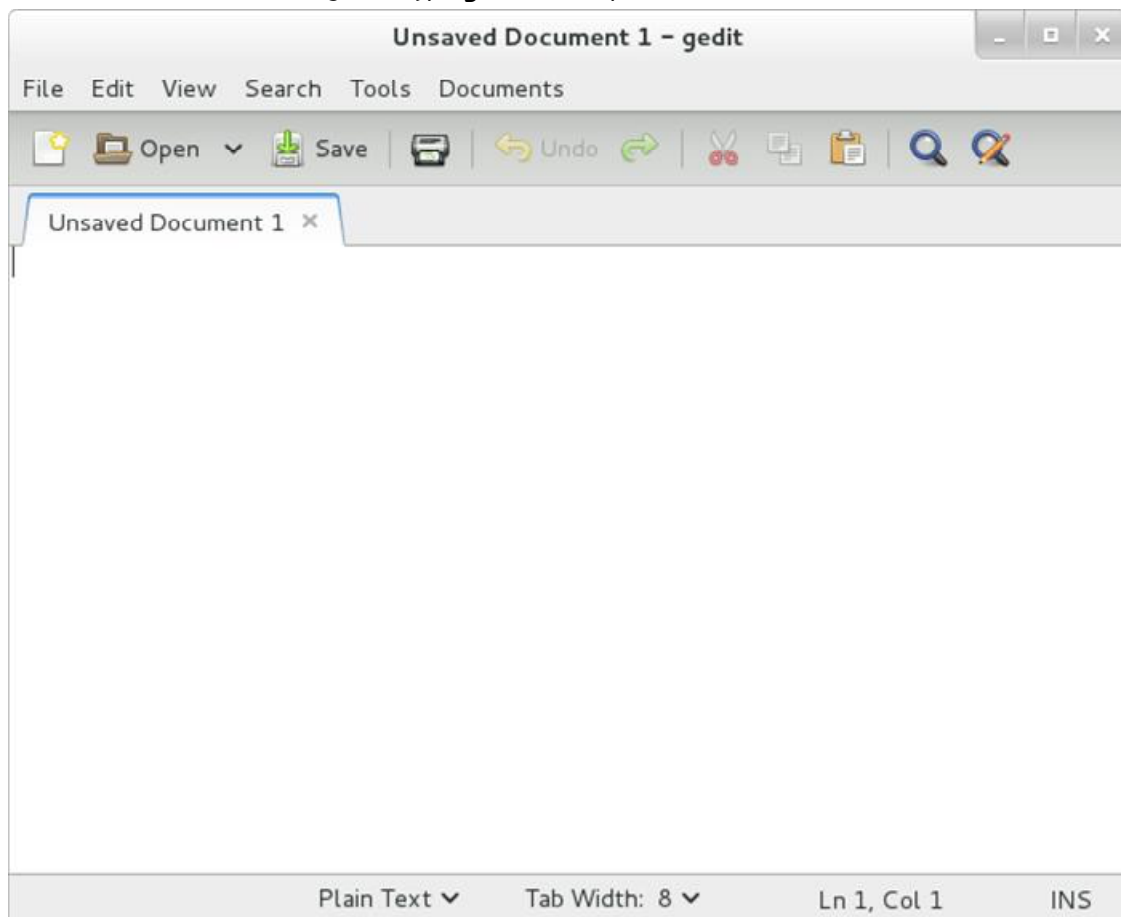
- Edit text files with `gedit`.
- Copy text between graphical windows.

### Editing files with `gedit`

*Editing text files with a graphical editor*

The `gedit` application is a full-featured text editor for the GNOME desktop environment.

Launch `gedit` by selecting **Applications** → **Accessories** → `gedit` from the GNOME menu. Like other graphical applications, `gedit` can be started without navigating the menu. Press **Alt+F2** to open the **Enter a Command** dialog box. Type `gedit` and press **Enter**.



*Figure 4.11: gedit text editor*

GNOME Help includes a `gedit` help guide, which may be found by selecting **Applications** → **Favorites** → **Help** from the GNOME menu. Then select **Go** → **All Documents** to view the list of graphical applications. Scroll down to select the `gedit` **Text Editor** hyperlink.

#### Basic `gedit` Keystrokes

Perform many file management tasks using `gedit`'s menu:

- To create a new file in gedit, click the **blank paper** toolbar icon, or select **File** → **New (Ctrl-n)** from the menu.
- To save a file, click the **disk-drive save** toolbar icon, or select **File** → **Save (Ctrl-s)** from the menu.
- To open an existing file, click the **Open** toolbar icon, or select **File** → **Open (Ctrl-o)** from the menu. The **Open Files** dialog window will display from which users can locate and select the file to open.

Multiple files may be opened simultaneously, each with a unique tab under the menu bar. Tabs display a file name after being saved the first time.

## ▮ Copying text between graphical windows

Text can be copied between documents, text windows, and command windows in the graphical environment. Selected text is duplicated using *copy and paste* or moved using *cut and paste*. Whether cut or copied, the text is held in memory for pasting into another location.

To select text:

- Click and hold the left mouse button before the first character desired.
- Drag the mouse over and down until all required text is in a single highlighted selection, then release the left button. Do not click again with the left button, as that deselects the text.

To paste the selection, multiple methods can accomplish the same result. In the first method:

- Click the right mouse button anywhere on the text area just selected.
- From the resulting context menu, select *either cut or copy*.
- Move the mouse to the window or document where the text is to be placed, click the left mouse button to position where the text should go, and click the right mouse button again, now choosing **paste**.

Here is a shorter mouse technique to practice:

- First, select the text.
- Hover the mouse over the destination window and click the center mouse button, just once, to paste the text at the cursor.

This last method can only copy, not cut. The original text remains selected and can be deleted. As with other methods, the text remains in memory and can be repeatedly pasted.



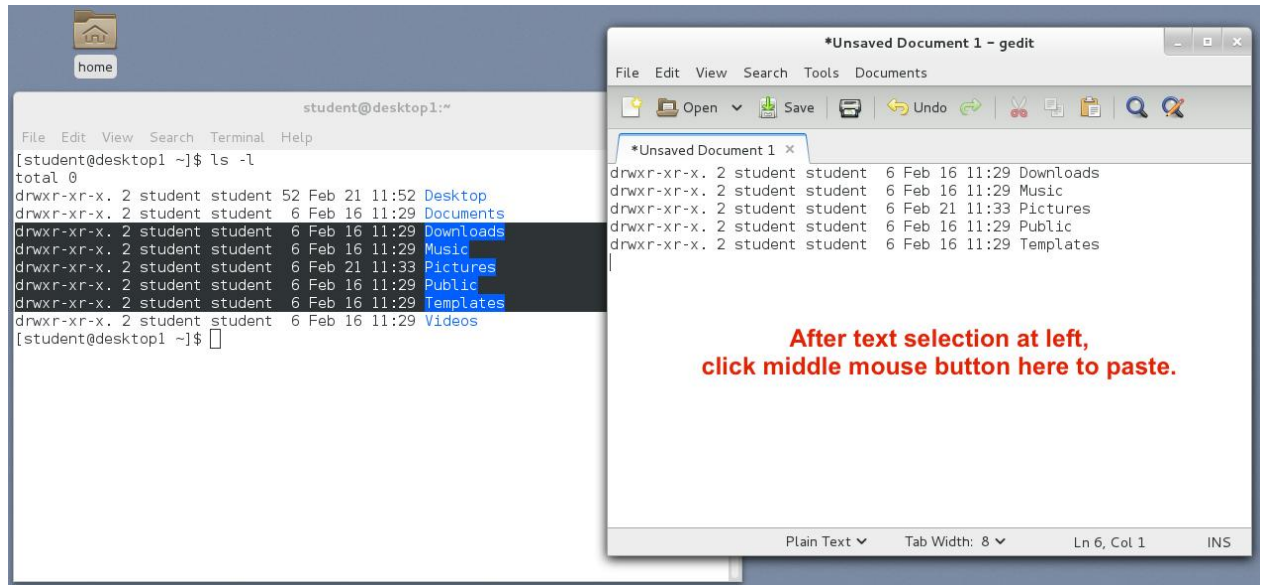


Figure 4.12: Select and paste using middle mouse button

The keyboard shortcut method can also be used in graphical applications:

- First, select the text.
- Use **Ctrl-x** to cut or **Ctrl-c** to copy the text.
- Click the location where the text is to be placed to position the cursor.
- Use **Ctrl-v** to paste.

#### Important

**Ctrl-c** and **Ctrl-v** will not copy and paste within a terminal window. **Ctrl-c** will actually terminate the current running process within a terminal window. To copy and paste within a terminal window, use **Ctrl-shift-c** and **Ctrl-shift-v**.

#### References

**gedit(1)** man page

gedit Text Editor

- **yelp help:gedit**

[gedit Wiki](#)

## 4.6. Practice: Copying Text Between Windows

### Guided exercise

In this lab, you will edit a file with gedit, selecting text and pasting it into the editor.

#### Outcomes:

An edited list of the configuration files found in the user's home directory.

#### Before you begin...

Perform the following steps on serverX unless directed otherwise. Log in as **student** and begin in student's home directory.

1. Redirect a long listing of all home directory files, including hidden, into a file named gedit\_lab.txt. Confirm that the file contains the listing.

2. `[student@serverX ~]$ cd`
3. `[student@serverX ~]$ ls -al > gedit_lab.txt`

```
[student@serverX ~]$ cat gedit_lab.txt
```

4. Open the file with the **gedit** text editor. Include the ending ampersand so that the shell prompt can return while **gedit** is running.

```
[student@serverX ~]$ gedit gedit_lab.txt &
```

5. Insert the date at the top of your file document.
  - a. In the shell command window, display today's date with day of the week, month, date, and year.

b. `[student@serverX ~]$ date +%A", "%B" "%d", "%Y`

```
Friday, February 21, 2014
```

- c. Select the text using the mouse.

```
[student@desktop1 ~]$
[student@desktop1 ~]$ date +%A", "%B" "%d", "%Y
Friday, February 21, 2014
[student@desktop1 ~]$
```

- d. Insert the text at the top of the file document. Switch to the **gedit** window. Using arrow keys, place the cursor at the upper-left corner of the document. Press the middle mouse button to paste the text.
  - e. Press **Enter** one or more times at the end of the inserted text to open blank lines above the file listing.
6. Insert a description for this document, including your username and host name, on line 2.
    - a. In the shell command window, create descriptive text using shell expansion concepts to include the username and host name where the file list was generated.

b. `[student@serverX ~]$ echo "$USER's  
configuration files on" $(hostname)`

```
student's configuration files on serverX.example.com
```

- c. Select the text using the mouse.

```
[student@desktop1 ~]$
[student@desktop1 ~]$ echo "$USER's configuration files on" $(hostname)
student's configuration files on desktop1.example.com
[student@desktop1 ~]$
```

- d. Insert the text on the second line of the file document. Switch to the **gedit** window. Using the arrow keys, place the cursor at the second line's leftmost character. Press the middle mouse button to paste the text.
  - e. Press **Enter** or **Delete**, as necessary, to maintain blank lines above the file listing.
7. Remove file lines that are not hidden configuration files or directories.
    - a. Remove the “total” line at the beginning of the listing.
    - b. Remove the two lines representing the current directory and the parent directory.
    - c. Remove lines for file names that do *not* start with a dot. Do not edit or remove lines for hidden files or directories that begin with a dot.
  8. The final file document should be similar to the following image. Manually edit the file to make corrections. The asterisk in the document tab or window header is a reminder of edits unsaved. Save the file and exit gedit.

The screenshot shows the gedit text editor window titled '\*gedit\_lab.txt (~) - gedit'. The menu bar includes File, Edit, View, Search, Tools, and Documents. The toolbar contains icons for Open, Save, Print, Undo, Redo, Cut, Copy, Paste, Find, and Replace. The document tab is labeled '\*gedit\_lab.txt x'. The text content is as follows:

```
Friday, February 21, 2014
student's configuration files on desktop1.example.com

-rw-----. 1 student student 194 Feb 21 16:00 .bash_history
-rw-r--r--. 1 student student 18 Aug 9 2013 .bash_logout
-rw-r--r--. 1 student student 193 Aug 9 2013 .bash_profile
-rw-r--r--. 1 student student 231 Aug 9 2013 .bashrc
drwx-----. 9 student student 4096 Feb 21 16:02 .cache
drwxr-xr-x. 15 student student 4096 Feb 21 16:11 .config
-rw-----. 1 student student 16 Feb 16 04:48 .esd_auth
-rw-----. 1 student student 1550 Feb 21 16:02 .ICEauthority
drwx-----. 3 student student 18 Feb 16 04:48 .local
drwx-----. 2 student student 28 Feb 15 13:47 .ssh
```

The status bar at the bottom indicates 'Plain Text', 'Tab Width: 8', 'Ln 12, Col 43', and 'INS'.

## 4.7. Lab: Creating, Viewing, and Editing Text Files

### Performance checklist

In this lab, you will edit a file using Vim's visual mode to simplify repetitive edits.

#### Outcomes:

Familiarity with the utilities and techniques required to perform file editing. The final edited file will be a list of selected files and tabular data.

#### Before you begin...

Perform the following steps on serverX unless directed otherwise. Log in as **student** and begin in the student's home directory.

1. Redirect a long listing of all content in student's home directory, including hidden directories and files, into a file named **editing\_final\_lab.txt**. Your home directory files may not exactly match those shown in the example graphics. This lab edits arbitrary lines and columns. The important outcome is to practice the visual selection process.

```
[student@serverX ~]$ cd
```

```
[student@serverX ~]$ ls -al > editing_final_lab.txt
```

3. Edit the file using Vim, to take advantage of *visual mode*.


```
[student@serverX ~]$ vim editing_final_lab.txt
```

4. Remove the first three lines, since those lines are not normal file names. Enter line-based visual mode with upper case **V**.

Use the arrow keys to position the cursor at the first character in the first row. Enter line-based visual mode with **V**. Move down using the down arrow key twice to select the first three rows. Delete the rows with **x**.

5. Remove the permission columns for group and world on the first line. In this step, enter visual mode with lower case **v**, which allows selecting characters on a single line only.

Use the arrow keys to position the cursor at the first character. Enter visual mode with **v**. Use the arrow keys to position the cursor at the last character, as shown in the screenshot. Delete the selection with **x**.




The screenshot shows a terminal window titled "student@desktop1:~" with a menu bar (File, Edit, View, Search, Terminal, Help). The terminal displays the output of the command `ls -al`, listing files and directories in the student's home directory. The first line of the listing is `-rw----- 1 student student 7691 Mar 5 10:56 .bash_history`. The cursor is positioned at the first character of this line. The terminal is in Vim's visual mode, indicated by the `-- VISUAL --` prompt at the bottom left. The status bar at the bottom right shows `1,11` and `All`.

```
student@desktop1:~
File Edit View Search Terminal Help
-rw----- 1 student student 7691 Mar 5 10:56 .bash_history
-rw-r--r-- 1 student student 18 Jan 29 05:45 .bash_logout
-rw-r--r-- 1 student student 193 Jan 29 05:45 .bash_profile
-rw-r--r-- 1 student student 231 Jan 29 05:45 .bashrc
drwx----- 12 student student 4096 Feb 22 13:23 .cache
drwxr-xr-x 18 student student 4096 Feb 21 11:33 .config
drwxr-xr-x 2 student student 6 Feb 21 20:06 Desktop
drwxr-xr-x 2 student student 4096 Feb 23 17:46 Documents
drwxr-xr-x 2 student student 6 Feb 16 11:29 Downloads
drwxr-xr-x 2 student student 4096 Feb 23 14:06 Music
drwxr-xr-x 2 student student 6 Feb 23 16:23 Pictures
drwxr-xr-x 2 student student 6 Feb 16 11:29 Public
drwx----- 2 student student 24 Feb 22 15:33 .ssh
drwxr-xr-x 2 student student 6 Feb 16 11:29 Templates
drwxr-xr-x 2 student student 4096 Feb 23 16:35 Videos
-rw----- 1 student student 1020 Feb 21 21:14 .viminfo
~
~
~
~
~
~
-- VISUAL -- 1,11 All
```

6. Remove the permission columns for group and world on the remaining lines. This step will use a more efficient block selection visual mode to avoid having to repeat the single line edit multiple times. This time, enter visual mode with the control sequence **Ctrl-v**, which allows selecting a block of characters on multiple lines.

Use the arrow keys to position the cursor at the first character. Enter visual mode with the control sequence **Ctrl-v**. Use the arrow keys to position the cursor at the last character of the column on the last line, as shown in the screenshot. Delete the selection with **x**.



The screenshot shows a terminal window titled "student@desktop1:~" with a menu bar (File, Edit, View, Search, Terminal, Help). The terminal displays a file listing with columns for permissions, file size, owner, group, size, date, time, and filename. A visual block selection is active, covering the permission and group owner columns of the first 15 lines. The selection is indicated by a dark grey background. At the bottom of the terminal, the status bar shows "-- VISUAL BLOCK --" on the left, "16,11" in the center, and "All" on the right.

```
student@desktop1:~
File Edit View Search Terminal Help
-rw- 1 student student 7691 Mar 5 10:56 .bash_history
-rw-r--r--. 1 student student 18 Jan 29 05:45 .bash_logout
-rw-r--r--. 1 student student 193 Jan 29 05:45 .bash_profile
-rw-r--r--. 1 student student 231 Jan 29 05:45 .bashrc
drwx----- 12 student student 4096 Feb 22 13:23 .cache
drwxr-xr-x. 18 student student 4096 Feb 21 11:33 .config
drwxr-xr-x. 2 student student 6 Feb 21 20:06 Desktop
drwxr-xr-x. 2 student student 4096 Feb 23 17:46 Documents
drwxr-xr-x. 2 student student 6 Feb 16 11:29 Downloads
drwxr-xr-x. 2 student student 4096 Feb 23 14:06 Music
drwxr-xr-x. 2 student student 6 Feb 23 16:23 Pictures
drwxr-xr-x. 2 student student 6 Feb 16 11:29 Public
drwx----- 2 student student 24 Feb 22 15:33 .ssh
drwxr-xr-x. 2 student student 6 Feb 16 11:29 Templates
drwxr-xr-x. 2 student student 4096 Feb 23 16:35 Videos
-rw----- 1 student student 1020 Feb 21 21:14 .viminfo
~
~
~
~
~
~
-- VISUAL BLOCK -- 16,11 All
```

7. Remove the *group owner* column, leaving only one "student" column on all lines. Use the same block selection technique as the last step.

Use the arrow keys to position the cursor at the first character of the group owner column. Enter visual mode with **Ctrl-v**. Use the arrow keys to position the cursor at the last character and row of the group owner column, as shown in the screenshot. Delete the selection with **x**.

```
student@desktop1:~
File Edit View Search Terminal Help
-rw- 1 student student 7691 Mar 5 10:56 .bash_history
-rw- 1 student student 18 Jan 29 05:45 .bash_logout
-rw- 1 student student 193 Jan 29 05:45 .bash_profile
-rw- 1 student student 231 Jan 29 05:45 .bashrc
drwx 12 student student 4096 Feb 22 13:23 .cache
drwx 18 student student 4096 Feb 21 11:33 .config
drwx 2 student student 6 Feb 21 20:06 Desktop
drwx 2 student student 4096 Feb 23 17:46 Documents
drwx 2 student student 6 Feb 16 11:29 Downloads
drwx 2 student student 4096 Feb 23 14:06 Music
drwx 2 student student 6 Feb 23 16:23 Pictures
drwx 2 student student 6 Feb 16 11:29 Public
drwx 2 student student 24 Feb 22 15:33 .ssh
drwx 2 student student 6 Feb 16 11:29 Templates
drwx 2 student student 4096 Feb 23 16:35 Videos
-rw- 1 student student 1020 Feb 21 21:14 .viminfo
~
~
~
~
~
~
-- VISUAL BLOCK -- 16,23 All
```

8. Remove the time column, but leave the month and day on all lines. Again, use the block selection visual mode.

Use the arrow keys to position the cursor at the first character. Enter visual mode with **Ctrl-v**. Use the arrow keys to position the cursor at the last character and row of the time column, as shown in the screenshot. Delete the selection with **x**.

```
student@desktop1:~
File Edit View Search Terminal Help
-rw- 1 student 7691 Mar 5 10:56 .bash_history
-rw- 1 student 18 Jan 29 05:45 .bash_logout
-rw- 1 student 193 Jan 29 05:45 .bash_profile
-rw- 1 student 231 Jan 29 05:45 .bashrc
drwx 12 student 4096 Feb 22 13:23 .cache
drwx 18 student 4096 Feb 21 11:33 .config
drwx 2 student 6 Feb 21 20:06 Desktop
drwx 2 student 4096 Feb 23 17:46 Documents
drwx 2 student 6 Feb 16 11:29 Downloads
drwx 2 student 4096 Feb 23 14:06 Music
drwx 2 student 6 Feb 23 16:23 Pictures
drwx 2 student 6 Feb 16 11:29 Public
drwx 2 student 24 Feb 22 15:33 .ssh
drwx 2 student 6 Feb 16 11:29 Templates
drwx 2 student 4096 Feb 23 16:35 Videos
-rw- 1 student 1020 Feb 21 21:14 .viminfo
~
~
~
~
~
~
-- VISUAL BLOCK -- 16,34 All
```

9. Remove the **Desktop** and **Public** rows. This time, enter visual mode with upper case **v**, which automatically selects full lines.

Use the arrow keys to position the cursor at any character on the **Desktop** row. Enter visual mode with upper case **V**. The full line is selected, as shown in the screenshot. Delete the selection with **x**. Repeat for the **Public** row.

```

student@desktop1:~
File Edit View Search Terminal Help
-rw- 1 student 7691 Mar 5 .bash_history
-rw- 1 student 18 Jan 29 .bash_logout
-rw- 1 student 193 Jan 29 .bash_profile
-rw- 1 student 231 Jan 29 .bashrc
drwx 12 student 4096 Feb 22 .cache
drwx 18 student 4096 Feb 21 .config
drwx 2 student 6 Feb 21 Desktop
drwx 2 student 4096 Feb 23 Documents
drwx 2 student 6 Feb 16 Downloads
drwx 2 student 4096 Feb 23 Music
drwx 2 student 6 Feb 23 Pictures
drwx 2 student 6 Feb 16 Public
drwx 2 student 24 Feb 22 .ssh
drwx 2 student 6 Feb 16 Templates
drwx 2 student 4096 Feb 23 Videos
-rw- 1 student 1020 Feb 21 .viminfo
~
~
~
~
~
-- VISUAL LINE --
7,31 All

```

10. Save and exit. Make a backup, using the date (in seconds) to create a unique file name.

```
[student@serverX ~]$ cp editing_final_lab.txt
editing_final_lab_$(date +%s).txt
```

11. Mail the file contents as the message, not an attachment, to the user student.

```
[student@serverX ~]$ cat editing_final_lab.txt | mail
-s "lab file" student
```

12. Append a dashed line to the file to recognize the beginning of newer content.

```
[student@serverX ~]$ echo
"-----" >>
editing_final_lab.txt
```

13. Append a full process listing, but only for processes owned by the current user **student** and running on the currently used terminal. View the process listing and send the listing to the file with one command line.

```
[student@serverX ~]$ ps -f | tee -a
editing_final_lab.txt
```

14. Confirm that the process listing is at the bottom of the lab file.

```
15. [student@serverX ~]$ cat editing_final_lab.txt
16. -rw- 1 student 7691 Mar 5 .bash_history
```



```
17. -rw- 1 student 18 Jan 29 .bash_logout
18. -rw- 1 student 193 Jan 29 .bash_profile
19. -rw- 1 student 231 Jan 29 .bashrc
20. drwx 12 student 4096 Feb 22 .cache
21. drwx 18 student 4096 Feb 21 .config
22. drwx 2 student 4096 Feb 23 Documents
23. drwx 2 student 6 Feb 16 Downloads
24. drwx 2 student 4096 Feb 23 Music
25. drwx 2 student 6 Feb 23 Pictures
26. drwx 2 student 24 Feb 22 .ssh
27. drwx 2 student 6 Feb 16 Templates
28. drwx 2 student 4096 Feb 23 Videos
29. -rw- 1 student 1020 Feb 21 .viminfo
30.
31. -----
32. UID PID PPID C STIME TTY TIME CMD
33. student 2005 2001 0 16:01 pts/0 00:00:00 /bin/bash
34. student 26923 2005 0 19:14 pts/0 00:00:00 ps -f
```

```
student 26924 2005 0 19:14 pts/0 00:00:00 tee -a editing_final_lab.txt
```

# Chapter 5. Managing Local Linux Users and Groups

## 5.1. Users and Groups

## 5.2. Practice: User and Group Concepts

## 5.3. Gaining Superuser Access

## 5.4. Practice: Running Commands as root

## 5.5. Managing Local User Accounts

## 5.6. Practice: Creating Users Using Command-line Tools

## 5.7. Managing Local Group Accounts

## 5.8. Practice: Managing Groups Using Command-line Tools

## 5.9. Managing User Passwords

## 5.10. Practice: Managing User Password Aging

## 5.11. Lab: Managing Local Linux Users and Groups

### Introduction

| Overview   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Goal       | To manage local Linux users and groups and administer local password policies.                                                                                                                                                                                                                                                                                                                                                                                                 |
| Objectives | <ul style="list-style-type: none"><li>○ Explain the role of users and groups on a Linux system and how they are understood by the computer.</li><li>○ Run commands as the superuser to administer a Linux system.</li><li>○ Create, modify, lock, and delete locally defined user accounts.</li><li>○ Create, modify, and delete locally defined group accounts.</li><li>○ Lock accounts manually or by setting a password-aging policy in the shadow password file.</li></ul> |
| Sections   | <ul style="list-style-type: none"><li>○ Users and Groups (and Practice)</li><li>○ Gaining Superuser Access (and Practice)</li><li>○ Managing Local User Accounts (and Practice)</li><li>○ Managing Local Group Accounts (and Practice)</li><li>○ Managing User Passwords (and Practice)</li></ul>                                                                                                                                                                              |
| Lab        | <ul style="list-style-type: none"><li>○ Managing Local Linux Users and Groups</li></ul>                                                                                                                                                                                                                                                                                                                                                                                        |

## 5.1. Users and Groups

### Objectives

After completing this section, students should be able to explain the role of users and groups on a Linux system and how they are understood by the computer.

### What is a user?

#### *Users and groups*

Every process (running program) on the system runs as a particular user. Every file is owned by a particular user. Access to files and directories are restricted by user. The user associated with a running process determines the files and directories accessible to that process.

The `id` command is used to show information about the current logged-in user. Basic information about another user can also be requested by passing in the username of that user as the first argument to the `id` command.

```
[student@desktopX ~]$ id
```

```
uid=1000(student) gid=1000(student) groups=1000(student),10(wheel)
context=unconfined_u:unconfined_r:unconfined_t:s0-s0:c0.c1023
```

To view the user associated with a file or directory, use the **ls -l** command. The third column shows the username:

```
[student@serverX ~]$ ls -l /tmp
drwx-----. 2 gdm gdm 4096 Jan 24 13:05 orbit-gdm
drwx-----. 2 student student 4096 Jan 25 20:40 orbit-student
-rw-r--r--. 1 root root 23574 Jan 24 13:05 postconf
```

To view process information, use the **ps** command. The default is to show only processes in the current shell. Add the **a** option to view all processes with a terminal. To view the user associated with a process, include the **u** option. The first column shows the username:

```
[student@serverX ~]$ ps au
USER PID %CPU %MEM VSZ RSS TTY STAT START TIME COMMAND
root 428 0.0 0.7 152768 14400 tty1 Ss+ Feb03 0:04 /usr/bin/Xorg
root 511 0.0 0.0 110012 812 ttyS0 Ss+ Feb03 0:00 /sbin/agetty
root 1805 0.0 0.1 116040 2580 pts/0 Ss Feb03 0:00 -bash
root 2109 0.0 0.1 178468 2200 pts/0 S Feb03 0:00 su - student
student 2110 0.0 0.1 116168 2864 pts/0 S Feb03 0:00 -bash
student 3690 0.0 0.0 123368 1300 pts/0 R+ 11:42 0:00 ps au
```

The output of the previous commands displays users by name, but internally the operating system tracks users by *aUID number*. The mapping of names to numbers is defined in databases of account information. By default, systems use a simple "flat file," the **/etc/passwd** file, to store information about local users. The format of **/etc/passwd** follows (seven colon-separated fields):

```
1 username: 2 password: 3 UID: 4 GID: 5 GECOS: 6 /home/dir: 7 shell
```

- 1 *username* is a mapping of a UID to a name for the benefit of human users.
- 2 *password* is where, historically, passwords were kept in an encrypted format. Today, they are stored in a separate file called **/etc/shadow**.
- 3 *UID* is a user ID, a number that identifies the user at the most fundamental level.
- 4 *GID* is the user's primary group ID number. Groups will be discussed in a moment.

- 5 *GECOS* field is arbitrary text, which usually includes the user's real name.
- 6 */home/dir* is the location of the user's personal data and configuration files.
- 7 *shell* is a program that runs as the user logs in. For a regular user, this is normally the program that provides the user's command line prompt.

## What is a group?

Like users, groups have a name and a number (GID). Local groups are defined in **/etc/group**.

### Primary groups

- Every user has exactly one *primary group*.
- For local users, the primary group is defined by the GID number of the group listed in the third field of **/etc/passwd**.
- Normally, the primary group owns new files created by the user.
- Normally, the primary group of a newly created user is a newly created group with the same name as the user. The user is the only member of this *User Private Group* (UPG).

### Supplementary groups

- Users may be a member of zero or more *supplementary groups*.
- The users that are supplementary members of local groups are listed in the last field of the group's entry in **/etc/group**. For local groups, user membership is determined by a comma-separated list of users found in the last field of the group's entry in **/etc/group**:

```
groupname:password:GID:list,of,users,in,this,group
```

- Supplementary group membership is used to help ensure that users have access permissions to files and other resources on the system.

### References

**id**(1), **passwd**(5), and **group**(5) man pages

**info libc** (*GNU C Library Reference Manual*)

- Section 29: Users and groups

(Note that the *glibc-devel* package must be installed for this info node to be available.)

## 5.2. Practice: User and Group Concepts

### Quiz

Match the items below to their counterparts in the table.

| Description                                                     | Keyword |
|-----------------------------------------------------------------|---------|
| A number that identifies the user at the most fundamental level | UID     |

| Description                                                      | Keyword                  |
|------------------------------------------------------------------|--------------------------|
| The program that provides the user's command line prompt         | login shell              |
| Location of local group information                              | <code>/etc/group</code>  |
| Location of the user's personal files                            | home directory           |
| A number that identifies the group at the most fundamental level | GID                      |
| Location of local user account information                       | <code>/etc/passwd</code> |
| The fourth field of <code>/etc/passwd</code>                     | primary group            |

## 5.3. Gaining Superuser Access

### Objectives

After completing this section, students should be able to run commands as the superuser to administer a Linux system.

### The `root` user

#### *Gaining superuser access*

Most operating systems have some sort of *superuser*, a user that has all power over the system. This user in Red Hat Enterprise Linux is the `root` user. This user has the power to override normal privileges on the file system, and is used to manage and administer the system. In order to perform tasks such as installing or removing software and to manage system files and directories, a user must escalate privileges to the `root` user.

Most devices can only be controlled by `root`, but there are a few exceptions. For instance, removable devices, such as USB devices, are allowed to be controlled by a normal user. Thus, a non-root user is allowed to add and remove files and otherwise manage a removable device, but only root is allowed to manage "fixed" hard drives by default.

This unlimited privilege, however, comes with responsibility. `root` has unlimited power to damage the system: remove files and directories, remove user accounts, add backdoors, etc. If the `root` account is compromised, someone else would have administrative control of the system. Throughout this course, administrators will be encouraged to log in as a normal user and escalate privileges to `root` only when needed.

The `root` account on Linux is roughly equivalent to the local Administrator account on Windows. In Linux, most system administrators log into an unprivileged user account and use various tools to temporarily gain root privileges.

#### Warning

One common practice on Windows in the past is for the local Administrator user to log in directly to perform system administrator duties. However, on Linux, it is recommended that system administrators *should not* log in directly as `root`. Instead, system administrators should log in as a non-root user, and use other mechanisms (`su`, `sudo`, or `PolicyKit`, for example) to temporarily gain superuser privileges.

By logging in as the administrative user, the entire desktop environment unnecessarily runs with administrative privileges. In that situation, any security vulnerability which would normally only compromise the user account has the potential to compromise the entire system.

In recent versions of Microsoft Windows, Administrator disabled by default, and features such as User Account Control (UAC) are used to limit administrative privileges for users until actually needed. In Linux, the **PolicyKit** system is the nearest equivalent to UAC.

## Switching users with `su`

The **su** command allows a user to switch to a different user account. If a username is not specified, the *root* account is implied. When invoked as a regular user, a prompt will display asking for the password of the account you are switching to; when invoked as *root*, there is no need to enter the account password.

```
su [-] <username>
```

```
[student@desktopX ~]$ su -
Password: redhat
[root@desktopX ~]#
```

The command **su *username*** starts a *non-login shell*, while the command **su - *username*** starts a *login shell*. The main distinction is **su -** sets up the shell environment as if this were a clean login as that user, while **su** just starts a shell as that user with the current environment settings.

In most cases, administrators want to run **su -** to get the user's normal settings. For more information, see the **bash(1)** man page.

### Note

The **su** command is most frequently used to get a command line interface (shell prompt) which is running as another user, typically *root*. However, with the **-c** option, it can be used like the Windows utility **runas** to run an arbitrary program as another user. See **info su** for details.

## Running commands as *root* with `sudo`

Fundamentally, Linux implements a very coarse-grained permissions model: *root* can do everything, other users can do nothing (systems-related). The common solution previously discussed is to allow standard users to temporarily “become *root*” using the **su** command. The disadvantage is that while acting as *root*, all the privileges (and responsibilities) of *root* are granted. Not only can the user restart the web server, but they can also remove the entire **/etc** directory. Additionally, all users requiring superuser privilege in this manner must know the *root* password.

The **sudo** command allows a user to be permitted to run a command as *root*, or as another user, based on settings in the **/etc/sudoers** file. Unlike other tools such as **su**, **sudo** requires users to enter their own password for authentication, not the password of the account they are trying to access. This allows an administrator to hand out fine-grained permissions to users to delegate system administration tasks, without having to hand out the *root* password.

For example, when **sudo** has been configured to allow the user *student* to run the command **usermod** as *root*, *student* could run the following command to lock a user account:

```
[student@serverX ~]$ sudo usermod -L username
```

```
[sudo] password for student: password
```

One additional benefit to using **sudo** is that all commands executed using **sudo** are logged by default to **/var/log/secure**.

```
[student@serverX ~]$ sudo tail /var/log/secure
...
Feb 19 15:23:36 localhost sudo: student : TTY=pts/0 ; PWD=/home/student ;
USER=root ; COMMAND=/sbin/usermod -L student
Feb 19 15:23:36 localhost usermod[16325]: lock user 'student' password
Feb 19 15:23:47 localhost sudo: student : TTY=pts/0 ; PWD=/home/student ;
USER=root ; COMMAND=/bin/tail /var/log/secure
```

In Red Hat Enterprise Linux 7, all members of group **wheel** can use **sudo** to run commands as any user, including **root**. The user will be prompted for their own password. This is a change from Red Hat Enterprise Linux 6 and earlier. Users who were members of group **wheel** did not get this administrative access by default in RHEL 6 and earlier.

To enable similar behavior on earlier versions of Red Hat Enterprise Linux, use **visudo** to edit the configuration file and uncomment the line allowing the group **wheel** to run all commands.

```
[root@desktopX ~]# cat /etc/sudoers
...Output omitted...
Allows people in group wheel to run all commands
%wheel ALL=(ALL) ALL

Same thing without a password
%wheel ALL=(ALL) NOPASSWD: ALL
...Output omitted...
```

### Warning

RHEL 6 did not grant group **wheel** any special privileges by default. Sites which have been using this group may be surprised when RHEL 7 automatically grants all members of **wheel** full **sudo** privileges. This could lead to unauthorized users getting superuser access to RHEL 7 systems.

Historically, membership in group **wheel** has been used by Unix-like systems to grant or control superuser access.

Most system administration applications with a GUI use **PolicyKit** to prompt users for authentication and to manage root access. In Red Hat Enterprise Linux 7, **PolicyKit** may also prompt members of group **wheel** for their own password in order to get **root** privileges when using graphical tools.

This is similar to the way in which they can use **sudo** to get those privileges at the shell prompt. **PolicyKit** grants these privileges based on its own configuration settings, separate from **sudo**. Advanced students may be interested in the **pkexec(1)** and **polkit(8)** man pages for details on how this system works, but it is beyond the scope of this course.

### References

**su**(1) and **sudo**(8) man pages

**info libc** (*GNU C Library Reference Manual*)

- o Section 29.2: The Persona of a Process

(Note that the *glibc-devel* package must be installed for this info node to be available.)

## 5.4. Practice: Running Commands as root

### Guided exercise

In this lab, you will practice running commands as **root**.

#### Outcomes

Use the **su** with and without login scripts to switch users. Use **sudo** to run commands with privilege.

#### Before you begin...

Reset your serverX system.

1. Log into the GNOME desktop on serverX as **student** with a password of **student**.
2. Open a window with a Bash prompt.  
Select **Applications** → **Utilities** → **Terminal**.
3. Explore characteristics of the current student login environment.
  - a. View the user and group information and display the current working directory.

```
b. [student@serverX ~]$ id
c. uid=1000(student) gid=1000(student) groups=1000(student),10(wheel)
 context=unconfined_u:unconfined_r:unconfined_t:s0-s0:c0.c1023
d. [student@serverX ~]$ pwd
e. /home/student
```

- f. View the variables which specify the home directory and the locations searched for executable files.

```
g. [student@serverX ~]$ echo $HOME
h. /home/student
i. [student@serverX ~]$ echo $PATH
j. /usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/home/student/.lo
 cal/bin:/home/student/bin
```

4. Switch to root without the dash and explore characteristics of the new environment.
  - a. Become the **root** user at the shell prompt.

```
b. [student@serverX ~]$ su
```

```
Password: redhat
```

- c. View the user and group information and display the current working directory. Note the identity changed, but not the current working directory.



```
d. [root@serverX student]# id
e. uid=0(root) gid=0(root) groups=0(root)
 context=unconfined_u:unconfined_r:unconfined_t:s0-s0:c0.c1023
f. [root@serverX student]# pwd
g. /home/student
```

- h. View the variables which specify the home directory and the locations searched for executable files. Look for references to the student and root accounts.

```
i. [root@serverX student]# echo $HOME
j. /root
k. [root@serverX student]# echo $PATH
l. /usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/home/student/.lo
 cal/bin:/home/student/bin
```

- m. Exit the shell to return to the **student** user.

```
n. [root@serverX student]# exit
o. exit
```

5. Switch to root with the dash and explore characteristics of the new environment.

- a. Become the **root** user at the shell prompt. Be sure all the login scripts are also executed.

```
b. [student@serverX ~]$ su -
```

```
Password: redhat
```

- c. View the user and group information and display the current working directory.

```
d. [root@serverX ~]# id
e. uid=0(root) gid=0(root) groups=0(root)
 context=unconfined_u:unconfined_r:unconfined_t:s0-s0:c0.c1023
f. [root@serverX ~]# pwd
g. /root
```

- h. View the variables which specify the home directory and the locations searched for executable files. Look for references to the student and root accounts.

```
i. [root@serverX ~]# echo $HOME
j. /root
k. [root@serverX ~]# echo $PATH
```

```
l. /usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin:/root/
bin
```

m. Exit the shell to return to the **student** user.

```
n. [root@serverX ~]# exit
o. logout
```

6. Run several commands as student which require root access.

a. View the last 5 lines of the **/var/log/messages**.

```
b. [student@serverX ~]$ tail -5 /var/log/messages
c. tail: cannot open '/var/log/messages' for reading: Permission denied
d. [student@serverX ~]$ sudo tail -5
/var/log/messages
e. Feb 3 15:07:22 localhost su: (to root) root on pts/0
f. Feb 3 15:10:01 localhost systemd: Starting Session 31 of user root.
g. Feb 3 15:10:01 localhost systemd: Started Session 31 of user root.
h. Feb 3 15:12:05 localhost su: (to root) root on pts/0
i. Feb 3 15:14:47 localhost su: (to student) root on pts/0
```

j. Make a backup of a configuration file in the **/etc** directory.

```
k. [student@serverX ~]$ cp /etc/motd /etc/motdOLD
l. cp: cannot create regular file '/etc/motdOLD': Permission denied
m. [student@serverX ~]$ sudo cp /etc/motd
/etc/motdOLD
```

n. Remove the **/etc/motdOLD** file that was just created.

```
o. [student@serverX ~]$ rm /etc/motdOLD
p. rm: remove write-protected regular empty file '/etc/motdOLD'? y
q. rm: cannot remove '/etc/motdOLD': Permission denied
r. [student@serverX ~]$ sudo rm /etc/motdOLD
```

s. Edit a configuration file in the **/etc** directory.

```
t. [student@serverX ~]$ echo "Welcome to class" >>
/etc/motd
u. -bash: /etc/motd: Permission denied
```

```
[student@serverX ~]$ sudo vim /etc/motd
```

## 5.5. Managing Local User Accounts

### Objectives

After completing this section, students should be able to create, modify, lock, and delete locally defined user accounts.

### Managing local users

*Managing local users*

A number of command-line tools can be used to manage local user accounts.

**useradd** creates users

- **useradd *username*** sets reasonable defaults for all fields in `/etc/passwd` when run without options. The **useradd** command does not set any valid password by default, and the user cannot log in until a password is set.
- **useradd --help** will display the basic options that can be used to override the defaults. In most cases, the same options can be used with the **usermod** command to modify an existing user.
- Some defaults, such as the range of valid UID numbers and default password aging rules, are read from the `/etc/login.defs` file. Values in this file are only used when creating new users. A change to this file will not have an effect on any existing users.

**usermod** modifies existing users

- **usermod --help** will display the basic options that can be used to modify an account. Some common options include:

| <b>usermod options:</b>          |                                                                                                                                     |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| <b>-c, --comment<br/>COMMENT</b> | Add a value, such as a full name, to the GECOS field.                                                                               |
| <b>-g, --gid GROUP</b>           | Specify the primary group for the user account.                                                                                     |
| <b>-G, --groups<br/>GROUPS</b>   | Specify a list of supplementary groups for the user account.                                                                        |
| <b>-a, --append</b>              | Used with the <b>-G</b> option to append the user to the supplemental groups mentioned without removing the user from other groups. |
| <b>-d, --home<br/>HOME_DIR</b>   | Specify a new home directory for the user account.                                                                                  |
| <b>-m, --move-home</b>           | Move a user home directory to a new location. Must be used with the <b>-d</b> option.                                               |

| usermod options:         |                                                 |
|--------------------------|-------------------------------------------------|
| <b>-s, --shell SHELL</b> | Specify a new login shell for the user account. |
| <b>-L, --lock</b>        | Lock a user account.                            |
| <b>-U, --unlock</b>      | Unlock a user account.                          |

### userdel deletes users

- **userdel *username*** removes the user from `/etc/passwd`, but leaves the home directory intact by default.
- **userdel -r *username*** removes the user and the user's home directory.

#### Warning

When a user is removed with **userdel** without the **-r** option specified, the system will have files that are owned by an unassigned user ID number. This can also happen when files created by a deleted user exist outside their home directory. This situation can lead to information leakage and other security issues.

In Red Hat Enterprise Linux 7 the **useradd** command assigns new users the first free UID number available in the range starting from UID 1000 or above. (unless one is explicitly specified with the **-u *UID*** option). This is how information leakage can occur: If the first free UID number had been previously assigned to a user account which has since been removed from the system, the old user's UID number will get reassigned to the new user, giving the new user ownership of the old user's remaining files. The following scenario demonstrates this situation:

```
[root@serverX ~]# useradd prince
[root@serverX ~]# ls -l /home
drwx-----. 3 prince prince 74 Feb 4 15:22 prince

[root@serverX ~]# userdel prince
[root@serverX ~]# ls -l /home
drwx-----. 3 1000 1000 74 Feb 4 15:22 prince

[root@serverX ~]# useradd bob
[root@serverX ~]# ls -l /home
drwx-----. 3 bob bob 74 Feb 4 15:23 bob
drwx-----. 3 bob bob 74 Feb 4 15:22 prince
```

Notice that **bob** now owns all files that **prince** once owned. Depending on the situation, one solution to this problem is to remove all "unowned" files from the system when the user that created them is deleted. Another solution is to manually assign the "unowned" files to a different user. The root user can find "unowned" files and directories by running: **find / -nouser -o -nogroup 2> /dev/null**.

### id displays user information

- **id** will display user information, including the user's UID number and group memberships.
- **id *username*** will display user information for *username*, including the user's UID number and group memberships.

### **passwd** sets passwords

- **passwd *username*** can be used to either set the user's initial password or change that user's password.
- The *root* user can set a password to any value. A message will be displayed if the password does not meet the minimum recommended criteria, but is followed by a prompt to retype the new password and all tokens are updated successfully.

```
[root@serverX ~]# passwd student
Changing password for user student.
New password: redhat123
BAD PASSWORD: The password fails the dictionary check - it is based on a dictionary
word
Retype new password: redhat123
```

```
passwd: all authentication tokens updated successfully.
```

- A regular user must choose a password which is at least 8 characters in length and is not based on a dictionary word, the username, or the previous password.

### **UID ranges**

Specific UID numbers and ranges of numbers are used for specific purposes by Red Hat Enterprise Linux.

- **UID 0** is always assigned to the superuser account, **root**.
- **UID 1-200** is a range of "system users" assigned statically to system processes by Red Hat.
- **UID 201-999** is a range of "system users" used by system processes that do not own files on the file system. They are typically assigned dynamically from the available pool when the software that needs them is installed. Programs run as these "unprivileged" system users in order to limit their access to just the resources they need to function.
- **UID 1000+** is the range available for assignment to regular users.

#### **Note**

Prior to Red Hat Enterprise Linux 7, the convention was that UID 1-499 was used for system users and UID 500+ for regular users. Default ranges used by **useradd** and **groupadd** can be changed in the **/etc/login.defs** file.

#### **References**

**useradd(8)**, **usermod(8)**, **userdel(8)** man pages

## 5.6. Practice: Creating Users Using Command-line Tools

### Guided exercise

In this lab, you will create a number of users on your serverX system, setting and recording an initial password for each user.

#### Outcomes

A system with additional user accounts.

#### Before you begin...

Reset your serverX system.

1. Log into the GNOME desktop on serverX as **student** with a password of **student**.
2. Open a window with a Bash prompt.  
Select **Applications** → **Utilities** → **Terminal**.
3. Become the **root** user at the shell prompt.

```
4. [student@serverX ~]$ su -
```

```
Password: redhat
```

5. Add the user *juliet*.

```
[root@serverX ~]# useradd juliet
```

6. Confirm that *juliet* has been added by examining the `/etc/passwd` file.

```
7. [root@serverX ~]# tail -2 /etc/passwd
```

```
8. tcpdump:x:72:72:::/sbin/nologin
```

```
9. juliet:x:1001:1001::/home/juliet:/bin/bash
```

10. Use the **passwd** command to initialize *juliet's* password.

```
11. [root@serverX ~]# passwd juliet
```

```
12. Changing password for user juliet.
```

```
13. New password: juliet
```

```
14. BAD PASSWORD: The password is shorter than 8 characters
```

```
15. Retype new password: juliet
```

```
passwd: all authentication tokens updated successfully.
```

16. Continue adding the remaining users in the steps below and set initial passwords.

- a. romeo

```
b. [root@serverX ~]# useradd romeo
```

```
c. [root@serverX ~]# passwd romeo
```

- d. Changing password for user romeo.
- e. New password: **romeo**
- f. BAD PASSWORD: The password is shorter than 8 characters
- g. Retype new password: **romeo**

```
passwd: all authentication tokens updated successfully.
```

h. hamlet

```
[root@serverX ~]# useradd hamlet
```

```
[root@serverX ~]# passwd hamlet
```

j. reba

```
[root@serverX ~]# useradd reba
```

```
[root@serverX ~]# passwd reba
```

l. dolly

```
[root@serverX ~]# useradd dolly
```

```
[root@serverX ~]# passwd dolly
```

n. elvis

```
[root@serverX ~]# useradd elvis
```

```
[root@serverX ~]# passwd elvis
```

## 5.7. Managing Local Group Accounts

### Objectives

After completing this section, students should be able to create, modify, and delete locally defined group accounts.

### Managing supplementary groups

*Managing local groups*

A group must exist before a user can be added to that group. Several command-line tools are used to manage local group accounts.

**groupadd** creates groups

- **groupadd** *groupname* without options uses the next available GID from the range specified in the `/etc/login.defs` file.
- The **-g** *GID* option is used to specify a specific GID.

```
[student@serverX ~]$ sudo groupadd -g 5000 ateam
```

#### Note

Given the automatic creation of user private groups (GID 1000+), it is generally recommended to set aside a range of GID numbers to be used for supplementary groups. A higher range will avoid a collision with a system group (GID 0-999).

- The **-r** option will create a system group using a GID from the range of valid system GID numbers listed in the `/etc/login.defs` file.

```
[student@serverX ~]$ sudo groupadd -r appusers
```

### groupmod modifies existing groups

- The **groupmod** command is used to change a group name to a GID mapping. The **-n** option is used to specify a new name.

```
[student@serverX ~]$ sudo groupmod -n javaapp appusers
```

- The **-g** option is used to specify a new GID.

```
[student@serverX ~]$ sudo groupmod -g 6000 ateam
```

### groupdel deletes a group

- The **groupdel** command will remove a group.

```
[student@serverX ~]$ sudo groupdel javaapp
```

- A group may not be removed if it is the primary group of any existing user. As with **userdel**, check all file systems to ensure that no files remain owned by the group.

### usermod alters group membership

- The membership of a group is controlled with user management. Change a user's primary group with **usermod -ggroupname**.

```
[student@serverX ~]$ sudo usermod -g student student
```

- Add a user to a supplementary group with **usermod -aG groupname username**.

```
[student@serverX ~]$ sudo usermod -aG wheel elvis
```

#### Important

The use of the **-a** option makes **usermod** function in "append" mode. Without it, the user would be removed from *all other* supplementary groups.



## References

`group(5)`, `groupadd(8)`, `groupdel(8)`, and `usermod(8)` man pages

## 5.8. Practice: Managing Groups Using Command-line Tools

### Guided exercise

In this lab, you will add users to newly created supplementary groups.

#### Outcomes

The **shakespeare** group consists of **juliet**, **romeo**, and **hamlet**.

The **artists** group contains **reba**, **dolly**, and **elvis**.

#### Before you begin...

Perform the following steps on serverX unless directed otherwise.

1. Become the **root** user at the shell prompt.

```
2. [student@serverX ~]$ su -
```

```
Password: redhat
```

3. Create a supplementary group called **shakespeare** with a group ID of **30000**.

```
[root@serverX ~]# groupadd -g 30000 shakespeare
```

4. Create a supplementary group called **artists**.

```
[root@serverX ~]# groupadd artists
```

5. Confirm that *shakespeare* and *artists* have been added by examining the `/etc/group` file.

```
6. [root@serverX ~]# tail -5 /etc/group
```

```
7. reba:x:1004:
```

```
8. dolly:x:1005:
```

```
9. elvis:x:1006:
```

```
10. shakespeare:x:30000:
```

```
11. artists:x:30001:
```

12. Add the *juliet* user to the *shakespeare* group as a supplementary group.

```
[root@serverX ~]# usermod -G shakespeare juliet
```

13. Confirm that *juliet* has been added using the `id` command.

```
14. [root@serverX ~]# id juliet
```

```
15. uid=1001(juliet) gid=1001(juliet) groups=1001(juliet),30000(shakespeare)
```

16. Continue adding the remaining users to groups as follows:

- a. Add *romeo* and *hamlet* to the *shakespeare* group.

```
[root@serverX ~]# usermod -G shakespeare romeo
```

```
[root@serverX ~]# usermod -G shakespeare hamlet
```

- c. Add *reba*, *dolly*, and *elvis* to the *artists* group.

```
[root@serverX ~]# usermod -G artists reba
```

```
[root@serverX ~]# usermod -G artists dolly
```

```
[root@serverX ~]# usermod -G artists elvis
```

- f. Verify the supplemental group memberships by examining the `/etc/group` file.

```
[root@serverX ~]# tail -5 /etc/group
```

```
h. reba:x:1004:
```

```
i. dolly:x:1005:
```

```
j. elvis:x:1006:
```

```
k. shakespeare:x:30000:juliet,romeo,hamlet
```

```
artists:x:30001:reba,dolly,elvis
```

## 5.9. Managing User Passwords

### Objectives

After completing this section, students should be able to lock accounts manually or by setting a password-aging policy in the shadow password file.

### Shadow passwords and password policy

#### Managing passwords

In the distant past, encrypted passwords were stored in the world-readable `/etc/passwd` file. This was thought to be reasonably secure until dictionary attacks on encrypted passwords became common. At that point, the encrypted passwords, or "password hashes," were moved to the more secure `/etc/shadow` file. This new file also allowed password aging and expiration features to be implemented.

There are three pieces of information stored in a modern password hash:

```
1gCjLa2/z$6Pu0EK0AzfCjxjv2hoLOB/
```

- 1: The hashing algorithm. The number 1 indicates an MD5 hash. The number 6 appears when a SHA-512 hash is used.
- gCjLa2/z: The salt used to encrypt the hash. This is originally chosen at random. The salt and the unencrypted password are combined and encrypted to create the encrypted password

hash. The use of a salt prevents two users with the same password from having identical entries in the `/etc/shadow` file.

3. `6Pu0EK0AzfCjxjv2hoLOB/`: The encrypted hash.

When a user tries to log in, the system looks up the entry for the user in `/etc/shadow`, combines the salt for the user with the unencrypted password that was typed in, and encrypts them using the hashing algorithm specified. If the result matches the encrypted hash, the user typed in the right password. If the result doesn't match the encrypted hash, the user typed in the wrong password and the login attempt fails. This method allows the system to determine if the user typed in the correct password without storing that password in a form usable for logging in.

#### Note

Red Hat Enterprise Linux 6 and 7 support two new strong password hashing algorithms, `SHA-256` (algorithm `5`) and `SHA-512` (algorithm `6`). Both the salt string and the encrypted hash are longer for these algorithms. The default algorithm used for password hashes can be changed by the `root` user by running the command `authconfig --passalgo` with one of the arguments `md5`, `sha256`, or `sha512`, as appropriate.

Red Hat Enterprise Linux 7 defaults to using SHA-512 encryption.

#### `/etc/shadow` format

The format of `/etc/shadow` follows (nine colon-separated fields):

```
1 name: 2 password: 3 lastchange: 4 minage: 5 maxage: 6
warning: 7 inactive: 8 expire: 9 blank
```

- 1 The login *name*. This must be a valid account name on the system.
- 2 The encrypted *password*. A password field which starts with a exclamation mark means that the password is locked.
- 3 The date of the last *password change*, represented as the number of days since 1970. 01. 01.
- 4 The *minimum* number of days before a password may be changed, where 0 means "no minimum age requirement."
- 5 The *maximum* number of days before a password must be changed.
- 6 The *warning* period that a password is about to expire. Represented in days, where 0 means "no warning given."

- 7 The number of days an account remains active after a password has expired. A user may still log into the system and change the password during this period. After the specified number of days, the account is locked, becoming *inactive*.
- 8 The account *expiration* date, represented as the number of days since 1970.01.01.
- 9 This *blank* field is reserved for future use.

## ▮ Password aging

The following diagram relates the relevant password-aging parameters, which can be adjusted using **chage** to implement a password-aging policy.

```
chage -m 0 -M 90 -W 7 -I 14 username
```

**chage -d 0 username** will force a password update on next login.

**chage -l username** will list a username's current settings.

**chage -E YYYY-MM-DD** will expire an account on a specific day.

#### Note

The **date** command can be used to calculate a date in the future.

```
[student@serverX ~]$ date -d "+45 days"
```

```
Sat Mar 22 11:47:06 EDT 2014
```

## Restricting access

With the **chage** command, an account expiration can be set. Once that date is reached, the user cannot log into the system interactively. The **usermod** command can "lock" an account with the **-L** option.

```
[student@serverX ~]$ sudo usermod -L elvis
```

```
[student@serverX ~]$ su - elvis
```

```
Password: elvis
```

```
su: Authentication failure
```

When a user has left the company, the administrator may lock and expire an account with a single **usermod** command. The date must be given as the number of days since 1970.01.01.

```
[student@serverX ~]$ sudo usermod -L -e 1 elvis
```

Locking the account prevents the user from authenticating with a password to the system. It is the recommended method of preventing access to an account by an employee who has left the company. If the employee returns, the account can later be unlocked with **usermod -U USERNAME**. If the account was also expired, be sure to also change the expiration date.

### The nologin shell

Sometimes a user needs an account with a password to authenticate to a system, but does not need an interactive shell on the system. For example, a mail server may require an account to store mail and a password for the user to authenticate with a mail client used to retrieve mail. That user does not need to log directly into the system.

A common solution to this situation is to set the user's login shell to **/sbin/nologin**. If the user attempts to log into the system directly, the **nologin** "shell" will simply close the connection.

```
[root@serverX ~]# usermod -s /sbin/nologin student
```

```
[root@serverX ~]# su - student
```

```
Last login: Tue Feb 4 18:40:30 EST 2014 on pts/0
```

```
This account is currently not available.
```

#### Important

Use of the **nologin** shell prevents interactive use of the system, but does not prevent all access. A user may still be able to authenticate and upload or retrieve files through applications such as web applications, file transfer programs, or mail readers.

#### References

**chage(1), usermod(8), shadow(5), crypt(3)** man pages

## 5.10. Practice: Managing User Password Aging

### Guided exercise

In this lab, you will set unique password policies for users.

#### Outcomes

The password for *romeo* must be changed when the user first logs into the system, every 90 days thereafter, and the account expires in 180 days.

#### Before you begin...

Perform the following steps on *serverX* unless directed otherwise.

1. Explore locking and unlocking accounts.

- a. Lock the **romeo** account.

```
[student@serverX ~]$ sudo usermod -L romeo
```

- b. Attempt to log in as **romeo**.

```
c. [student@serverX ~]$ su - romeo
```

```
d. Password: romeo
```

```
e. su: Authentication failure
```

- f. Unlock the **romeo** account.

```
[student@serverX ~]$ sudo usermod -U romeo
```

2. Change the password policy for **romeo** to require a new password every 90 days.

```
3. [student@serverX ~]$ sudo chage -M 90 romeo
```

```
4. [student@serverX ~]$ sudo chage -l romeo
```

```
5. Last password change : Feb 03, 2014
```

```
6. Password expires : May 04, 2014
```

```
7. Password inactive : never
```

```
8. Account expires : never
```

```
9. Minimum number of days between password change : 0
```

```
10. Maximum number of days between password change : 90
```

```
11. Number of days of warning before password expires : 7
```

12. Additionally, force a password change on the first login for the **romeo** account.

```
13. [student@serverX ~]$ sudo chage -d 0 romeo
```

14. Log in as **romeo** and change the password to **forsooth123**.

```
15. [student@serverX ~]$ su - romeo
```

```

16. Password: romeo
17. You are required to change your password immediately (root enforced)
18. Changing password for romeo.
19. (current) UNIX password: romeo
20. New password: forsooth123
21. Retype new password: forsooth123
22. [romeo@serverX ~]$ exit

```

23. Expire accounts in the future.

a. Determine a date 180 days in the future.

```

b. [student@serverX ~]$ date -d "+180 days"
c. Sat Aug 2 17:05:20 EDT 2014

```

d. Set accounts to expire on that date.

```

e. [student@serverX ~]$ sudo chage -E 2014-08-02
 romeo
f. [student@serverX ~]$ sudo chage -l romeo
g. Last password change : Feb 03, 2014
h. Password expires : May 04, 2014
i. Password inactive : never
j. Account expires : Aug 02, 2014
k. Minimum number of days between password change : 0
l. Maximum number of days between password change : 90

```

```

Number of days of warning before password expires : 7

```

## 5.11. Lab: Managing Local Linux Users and Groups

### Performance checklist

In this lab, you will define a default password policy, create a supplementary group of three new users, and modify the password policy of one user.

#### Outcomes

- A new group on serverX called **consultants**, including three new user accounts for Sam Spade, Betty Boop, and Dick Tracy.
- All new accounts should require that passwords be changed at first login and every 30 days thereafter.
- The new consultant accounts should expire at the end of the 90-day contract, and Betty Boop must change her password every 15 days.

Before you begin...

Reset your serverX system.

1. Ensure that newly created users have passwords which must be changed every 30 days.

```
2. [student@serverX ~]$ sudo vim /etc/login.defs
3. [student@serverX ~]$ cat /etc/login.defs
4. ...Output omitted...
5. PASS_MAX_DAYS 30
6. PASS_MIN_DAYS 0
7. PASS_MIN_LEN 5
8. PASS_WARN_AGE 7
9. ...Output omitted...
```

10. Create a new group named **consultants** with a GID of 40000.

```
11. [student@serverX ~]$ sudo groupadd -g 40000
 consultants
12. [student@serverX ~]$ tail -5 /etc/group
13. stapdev:x:158:
14. pesign:x:989:
15. tcpdump:x:72:
16. slocate:x:21:
17. consultants:x:40000:
```

18. Create three new users: **sspade**, **bboop**, and **dtracy**, with a password of **default** and add them to the supplementary group **consultants**. The primary group should remain as the user private group.

```
19. [student@serverX ~]$ sudo useradd -G consultants
 sspade
20. [student@serverX ~]$ sudo useradd -G consultants bboop
21. [student@serverX ~]$ sudo useradd -G consultants
 dtracy
22. [student@serverX ~]$ tail -5 /etc/group
23. slocate:x:21:
24. consultants:x:40000:sspade,bboop,dtracy
25. sspade:x:1001:
26. bboop:x:1002:
27. dtracy:x:1003:
28. [student@serverX ~]$ sudo passwd sspade
29. Changing password for user sspade.
30. New password: default
```



```
31. BAD PASSWORD: The password is shorter than 8 characters
32. Retype new password: default
33. passwd: all authentication tokens updated successfully.
34. [student@serverX ~]$ sudo passwd bboop
35. [student@serverX ~]$ sudo passwd dtracy
```

36. Determine the date 90 days in the future and set each of the three new user accounts to expire on that date.

```
37. [student@serverX ~]$ date -d "+90 days"
38. Mon May 5 11:49:24 EDT 2014
39. [student@serverX ~]$ sudo chage -E 2014-05-05 sspade
40. [student@serverX ~]$ sudo chage -E 2014-05-05 bboop
41. [student@serverX ~]$ sudo chage -E 2014-05-05 dtracy
```

42. Change the password policy for the **bboop** account to require a new password every 15 days.

```
43. [student@serverX ~]$ sudo chage -M 15 bboop
44. [student@serverX ~]$ sudo chage -l bboop
45. Last password change : Feb 04, 2014
46. Password expires : Feb 19, 2014
47. Password inactive : never
48. Account expires : May 05, 2014
49. Minimum number of days between password change : 0
50. Maximum number of days between password change : 15
51. Number of days of warning before password expires : 7
```

52. Additionally, force all users to change their password on first login.

```
53. [student@serverX ~]$ sudo chage -d 0 sspade
54. [student@serverX ~]$ sudo chage -d 0 bboop
55. [student@serverX ~]$ sudo chage -d 0 dtracy
```

56. When you finish, run the **lab localusers grade** evaluation script to confirm you have done everything correctly.

```
[student@serverX ~]$ lab localusers grade
```

# Chapter 6. Controlling Access to Files with Linux File System Permissions

6.1. Linux File System Permissions

6.2. Practice: Interpreting File and Directory Permissions

6.3. Managing File System Permissions from the Command Line

6.4. Practice: Managing File Security from the Command Line

6.5. Managing Default Permissions and File Access

6.6. Practice: Controlling New File Permissions and Ownership

6.7. Lab: Controlling Access to Files with Linux File System Permissions

## Introduction

| Overview          |                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Goal</b>       | To set Linux file system permissions on files and interpret the security effects of different permission settings.                                                                                                                                                                                                                                                                      |
| <b>Objectives</b> | <ul style="list-style-type: none"><li>○ Explain how the Linux file permissions model works.</li><li>○ Change the permissions and ownership of files using command-line tools.</li><li>○ Configure a directory in which newly created files are automatically writable by members of the group which owns the directory, using special permissions and default umask settings.</li></ul> |
| <b>Sections</b>   | <ul style="list-style-type: none"><li>○ Linux File System Permissions (and Practice)</li><li>○ Managing File System Permissions from the Command Line (and Practice)</li><li>○ Managing Default Permissions and File Access (and Practice)</li></ul>                                                                                                                                    |
| <b>Lab</b>        | <ul style="list-style-type: none"><li>○ Controlling Access to Files with Linux File System Permissions</li></ul>                                                                                                                                                                                                                                                                        |

## 6.1. Linux File System Permissions

### Objectives

After completing this section, students should be able to explain how the Linux file permissions model works.

### Linux file system permissions

#### Linux file system permissions

Access to files by users are controlled by *file permissions*. The Linux file permissions system is simple but flexible, which makes it easy to understand and apply, yet able to handle most normal permission cases easily.

Files have just three categories of user to which permissions apply. The file is owned by a *user*, normally the one who created the file. The file is also owned by a single *group*, usually the primary group of the user who created the file, but this can be changed. Different permissions can be set for the owning user, the owning group, and for all *other* users on the system that are not the user or a member of the owning group.

The most specific permissions apply. So, *user* permissions override *group* permissions, which override *other* permissions.

In the graphic that follows, joshua is a member of the groups joshua and web, while allison is a member of allison, wheel, and web. When joshua and allison have the need to collaborate, the files should be associated with the group web and the group permissions should allow the desired access.

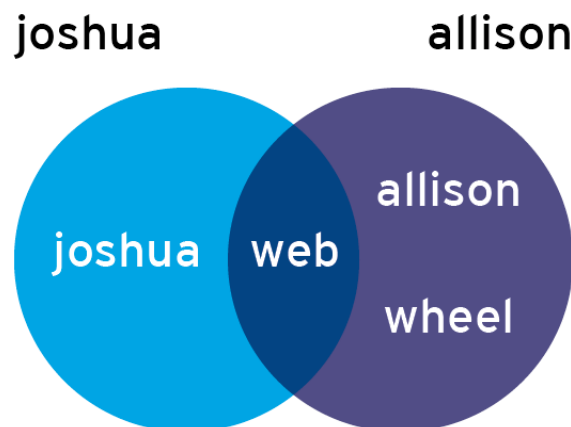


Figure 6.1: Group membership illustration

There are also just three categories of permissions which apply: *read*, *write*, and *execute*. These permissions affect access to files and directories as follows:

□

#### Effects of permissions on files and directories

| Permission       | Effect on files                      | Effect on directories                                 |
|------------------|--------------------------------------|-------------------------------------------------------|
| <b>r</b> (read)  | Contents of the file can be read.    | Contents of the directory (file names) can be listed. |
| <b>w</b> (write) | Contents of the file can be changed. | Any file in the directory may be created or deleted.  |

| Permission      | Effect on files                    | Effect on directories                                                                                   |
|-----------------|------------------------------------|---------------------------------------------------------------------------------------------------------|
| <b>x</b> (exec) | Files can be executed as commands. | Contents of the directory can be accessed (dependent on the permissions of the files in the directory). |

Note that users normally have both **read** and **exec** on read-only directories, so that they can list the directory and access its contents. If a user only has **read** access on a directory, the names of the files in it can be listed, but no other information, including permissions or time stamps, are available, nor can they be accessed. If a user only has **exec** access on a directory, they cannot list the names of the files in the directory, but if they already know the name of a file which they have permission to read, then they can access the contents of that file by explicitly specifying the file name.

A file may be removed by anyone who has write permission to the directory in which the file resides, regardless of the ownership or permissions on the file itself. (This can be overridden with a special permission, the *sticky bit*, which will be discussed at the end of the unit.)

### ▣ Viewing file/directory permissions and ownership

The **-l** option of the **ls** command will expand the file listing to include both the permissions of a file and the ownership:

```
[student@desktopX ~]$ ls -l test
-rw-rw-r--. 1 student student 0 Feb 8 17:36 test
```

The command **ls -l directoryname** will show the expanded listing of all of the files that reside inside the directory. To prevent the descent into the directory and see the expanded listing of the directory itself, add the **-ld** option to **ls**:

```
[student@desktopX ~]$ ls -ld /home
drwxr-xr-x. 5 root root 4096 Jan 31 22:00 /home
```

#### Note

Unlike NTFS permissions, Linux permissions only apply to the directory or file that they are set on. Permissions on a directory are not inherited automatically by the subdirectories and files within it. (The permissions on a directory *may* effectively block access to its contents, however.) All permissions in Linux are set directly on each file or directory.

The **read** permission on a directory in Linux is roughly equivalent to **List folder contents** in Windows.

The **write** permission on a directory in Linux is equivalent to **Modify** in Windows; it implies the ability to delete files and subdirectories. In Linux, if **write** and the **sticky bit** are both set on a directory, then only the user that owns a file or subdirectory in the directory may delete it, which is close to the behavior of the Windows **write** permission.

Root has the equivalent of the Windows **Full Control** permission on all files in Linux. However, root may still have access restricted by the system's SELinux policy and the security context of the process and files in question. SELinux will be discussed in a later course.

### ▣ Examples: Linux user, group, other concepts

```
Users and their groups:
```

```

lucy lucy,ricardo
ricky ricky,ricardo
ethel ethel,mertz
fred fred,mertz

```

File attributes (permissions, user & group ownership, name):

```

drwxrwxr-x ricky ricardo dir (which contains the following files)
-rw-rw-r-- lucy lucy lfile1
-rw-r--rw- lucy ricardo lfile2
-rw-rw-r-- ricky ricardo rfile1
-rw-r----- ricky ricardo rfile2

```

| Allowed/denied behavior                                                                                 | Controlling permissions                                                                                                                                                                               |
|---------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>lucy</b> is the only person who can change the contents of <b>lfile1</b> .                           | <b>lucy</b> has write permissions on the file <b>lfile1</b> as the owner. No one is listed as a member of the group <b>lucy</b> . The permissions for <i>other</i> do not include write permissions.  |
| <b>ricky</b> can view the contents of <b>lfile2</b> , but cannot modify the contents of <b>lfile2</b> . | <b>ricky</b> is a member of the group <b>ricardo</b> , and that group has read-only permissions on <b>lfile2</b> . Even though <i>other</i> has write permissions, group permissions take precedence. |
| <b>ricky</b> can delete <b>lfile1</b> and <b>lfile2</b> .                                               | <b>ricky</b> has write permissions on the directory containing both files, and as such, he can delete any file in that directory.                                                                     |
| <b>ethel</b> can change the contents of <b>lfile2</b> .                                                 | Since <b>ethel</b> is not <b>lucy</b> , and is not a member of the <b>ricardo</b> group, <i>other</i> permissions apply to her, and those include write permission.                                   |
| <b>lucy</b> can change the contents of <b>rfile1</b> .                                                  | <b>lucy</b> is a member of the <b>ricardo</b> group, and that group has both read and write permissions                                                                                               |

| Allowed/denied behavior                                                                                                                                                                                                                       | Controlling permissions                                                                                                                                                                                                                                                                                                                        |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                                                                                                                                                                                               | on <b>rfile1</b> .                                                                                                                                                                                                                                                                                                                             |
| <p><b>ricky</b> can view and modify the contents of <b>rfile2</b>.</p> <p><b>lucy</b> can view but not modify the contents of <b>rfile2</b>.</p> <p><b>ethel</b> and <b>fred</b> do not have any access to the contents of <b>rfile2</b>.</p> | <p><b>ricky</b> owns the file and has both read and write access to <b>rfile2</b>.</p> <p><b>lucy</b> is a member of the <b>ricardo</b> group, and that group has read-only access to <b>rfile2</b>.</p> <p><i>other</i> permissions apply to <b>ethel</b> and <b>fred</b>, and those permissions do not include read or write permission.</p> |

#### References

**ls(1)** man page

**info coreutils** (*GNU Coreutils*)

- o Section 13: Changing file attributes

## 6.2. Practice: Interpreting File and Directory Permissions

### Quiz

Using the directory listing presented, match the items that follow to their counterparts in the table.

Users and their groups:

```
wilma wilma,flintstone
fred fred,flintstone
betty betty,rubble
barney barney,rubble
```

File attributes (permissions, user & group ownership, name):

```
drwxrwxr-x fred flintstone dir (which contains the following files)
-rw-rw-r-- wilma wilma lfile1
-rw-r--rw- wilma flintstone lfile2
-rw-rw-r-- fred flintstone rfile1
-rw-r----- fred flintstone rfile2
```

| Description                                       | File name |
|---------------------------------------------------|-----------|
| Is owned by <b>fred</b> and readable by all users | rfile1    |
| Contents may be modified by the user <b>betty</b> | lfile2    |
| Can be deleted by the user <b>fred</b>            | all       |
| Cannot be read by the user <b>barney</b>          | rfile2    |
| Has a group ownership of <b>wilma</b>             | lfile1    |
| Can be deleted by the user <b>barney</b>          | none      |

## 6.3. Managing File System Permissions from the Command

### Line

#### ▮ Objectives

After completing this section, students should be able to change the permissions and ownership of files using command-line tools.

*Managing file system permissions from the command line*

#### ▮ Changing file/directory permissions

The command used to change permissions from the command line is **chmod**, short for "change mode" (permissions are also called the *mode* of a file). The **chmod** command takes a permission instruction followed by a list of files or directories to change. The permission instruction can be issued either symbolically (the symbolic method) or numerically (the numeric method).

**Symbolic method keywords:**

```
chmod WhoWhatWhich file|directory
```

- *Who* is u, g, o, a (for user, group, other, all)
- *What* is +, -, = (for add, remove, set exactly)
- *Which* is r, w, x (for read, write, executable)

The *symbolic* method of changing file permissions uses letters to represent the different groups of permissions: **u** for user, **g** for group, **o** for other, and **a** for all.

With the symbolic method, it is not necessary to set a complete new group of permissions. Instead, it is possible to change one or more of the existing permissions. In order to accomplish this, use three symbols: **+** to add permissions to a set, **-** to remove permissions from a set, and **=** to replace the entire set for a group of permissions.

The permissions themselves are represented by a single letter: **r** for read, **w** for write, and **x** for execute.

**Numeric method:**

```
chmod ### file|directory
```

- Each digit represents an access level: user, group, other.
- # is sum of r=4, w=2, and x=1.

Using the *numeric* method, permissions are represented by a three-digit (or four, when setting advanced permissions) *octal* number. A single octal digit can represent the numbers **0–7**, exactly the number of possibilities for a three-bit number.

To convert between symbolic and numeric representation of permissions, we need to know how the mapping is done. In the three-digit octal (numeric) representation, each digit stands for one group of permissions, from left to right: user, group, and other. In each of these groups, start with **0**. If the read permission is present, add **4**. Add **2** if write is present, and **1** for execute.

Numeric permissions are often used by advanced administrators since they are shorter to type and pronounce, while still giving full control over all permissions.

Examine the permissions **-rwxr-x---**. For the user, **rwx** is calculated as **4+2+1=7**. For the group, **r-x** is calculated as **4+0+1=5**, and for other users, **---** is represented with **0**. Putting these three together, the numeric representation of those permissions is **750**.

This calculation can also be performed in the opposite direction. Look at the permissions **640**. For the user permissions, **6** represents read (4) and write (2), which displays as **rw-**. For the group part, **4** only includes read (4) and displays as **r--**. The **0** for other provides no permissions (**---**) and the final set of symbolic permissions for this file is **-rw-r-----**.

**Examples**

- Remove read and write permission for group and other on **file1**:

```
[student@desktopX ~]$ chmod go-rw file1
```

- Add execute permission for everyone on **file2**:

```
[student@desktopX ~]$ chmod a+x file2
```

- Set read, write, and execute permission for user, read, and execute for group, and no permission for other on **sampledir**:

```
[student@desktopX ~]$ chmod 750 sampledir
```

#### Note

The **chmod** command supports the **-R** option for recursively setting permissions on an entire directory tree. When using this option, be sure to use the **X** permissions instead of the **x** permission to indicate that execute permissions should only be set on directories, and not regular files. For example, the following command will recursively set read and write access on **demodir** and all its children for their group owner, but will only apply execute permissions to directories, not regular files:



```
[student@desktopX ~]# chmod -R g+rwX demodir
```

## Changing file/directory user or group ownership

A newly created file is owned by the user who creates the file. By default, the new file has a group ownership which is the primary group of the user creating the file. Since Red Hat Enterprise Linux uses user private groups, this group is often a group with only that user as a member. To grant access based on group membership, the owner or the group of a file may need to be changed.

File ownership can be changed with the **chown** command (change owner). For example, to grant ownership of the file **foofile** to **student**, the following command could be used:

```
[root@desktopX ~]# chown student foofile
```

**chown** can be used with the **-R** option to recursively change the ownership of an entire directory tree. The following command would grant ownership of **foodir** and all files and subdirectories within it to **student**:

```
[root@desktopX ~]# chown -R student foodir
```

The **chown** command can also be used to change group ownership of a file by preceding the group name with a colon (:). For example, the following command will change the group **foodir** to **admins**:

```
[root@desktopX ~]# chown :admins foodir
```

The **chown** command can also be used to change both owner and group at the same time by using the syntax **owner:group**. For example, to change the ownership of **foodir** to **visitor** and the group to **guests**, use:

```
[root@desktopX ~]# chown visitor:guests foodir
```

Only **root** can change the ownership of a file. Group ownership, however, can be set by **root** or the file's owner. **root** can grant ownership to any group, while non-**root** users can grant ownership only to groups they belong to.

### Note

Instead of using **chown**, some users change the group ownership by using the **chgrp** command; this command works exactly the same as changing ownership with **chown**, including the use of **-R** to affect entire directory trees.

### References

**ls**(1), **chmod**(1), **chown**(1), and **chgrp**(1) man pages

## 6.4. Practice: Managing File Security from the Command

### Line

#### Guided exercise

In this lab, you will create a collaborative directory for pre-existing users.

#### Outcomes

A directory accessible by all members of the **ateam** group and a file created by Andy that can be modified by Alice.

### Before you begin...

Reset your serverX system.

1. Log into the GNOME desktop on serverX as **student** with a password of **password**.
2. Open a window with a Bash prompt.  
Select **Applications** → **Utilities** → **Terminal**.
3. Become the **root** user at the shell prompt.

```
4. [student@serverX ~]$ su -
```

```
Password: redhat
```

5. Run **lab permissions setup** which will create a shared group, **ateam**, with two new users, **andy** and **alice**. The password for these accounts is **password**

```
[root@serverX ~]# lab permissions setup
```

6. Create a directory in **/home** called **ateam-text**.

```
[root@serverX ~]# mkdir /home/ateam-text
```

7. Change the group ownership of the **ateam-text** directory to **ateam**.

```
[root@serverX ~]# chown :ateam /home/ateam-text
```

8. Ensure the permissions of **ateam-text** allows group members to create and delete files.

```
[root@serverX ~]# chmod g+w /home/ateam-text
```

9. Ensure the permissions of **ateam-text** forbids others from accessing its files.

```
10. [root@serverX ~]# chmod 770 /home/ateam-text
```

```
11. [root@serverX ~]$ ls -ld /home/ateam-text
```

```
drwxrwx---. 1 root ateam 6 Jan 23 12:50 /home/ateam-text
```

12. Exit the root shell and switch to the user **andy** with a password of **password**.

```
13. [root@serverX ~]# exit
```

```
14. [student@serverX ~]$ su - andy
```

```
Password: password
```

15. Navigate to the **/home/ateam-text** folder (remember to open a terminal window first).

```
[andy@serverX ~]$ cd /home/ateam-text
```

16. Create an empty file called **andyfile3**.

```
[andy@serverX ateam-text]$ touch andyfile3
```

17. Record the default user and group ownership of the new file and its permissions.

```
18. [andy@serverX ateam-text]$ ls -l andyfile3
```

```
-rw-rw-r--. 1 andy andy 0 Jan 23 12:59 andyfile3
```

19. Change the group ownership of the new file to **ateam** and record the new ownership and permissions.

```
20. [andy@serverX ateam-text]$ chown :ateam andyfile3
```

```
21. [andy@serverX ateam-text]$ ls -l andyfile3
```

```
-rw-rw-r--. 1 andy ateam 0 Jan 23 12:59 andyfile3
```

22. Exit the shell and switch to the user **alice** with a password of **password**.

```
23. [root@serverX ~]# exit
```

```
24. [student@serverX ~]$ su - alice
```

```
Password: password
```

25. Navigate to the **/home/ateam-text** folder.

```
[alice@serverX ~]$ cd /home/ateam-text
```

26. Determine **alice**'s privileges to access and/or modify **andyfile3**.

```
27. [alice@serverX ateam-text]$ echo "text" >> andyfile3
```

```
28. [alice@serverX ateam-text]$ cat andyfile3
```

```
text
```

## ▮ 6.5. Managing Default Permissions and File Access

### ▮ Objectives

After completing this section, students should be able to configure a directory in which newly created files are automatically writable by members of the group which owns the directory, using special permissions and default umask settings.

*Controlling default permissions on new files*

## Special permissions

The **setuid** (or **setgid**) permission on an executable file means that the command will run as the **user** (or **group**) of the file, not as the user that ran the command. One example is the **passwd** command:

```
[student@desktopX ~]$ ls -l /usr/bin/passwd
-rwsr-xr-x. 1 root root 35504 Jul 16 2010 /usr/bin/passwd
```

In a long listing, you can spot the **setuid** permissions by a lowercase **s** where you would normally expect the **x**(owner execute permissions) to be. If the owner does not have execute permissions, this will be replaced by an uppercase **S**.

The **sticky bit** for a directory sets a special restriction on deletion of files: Only the owner of the file (and **root**) can delete files within the directory. An example is **/tmp**:

```
[student@desktopX ~]$ ls -ld /tmp
drwxrwxrwt. 39 root root 4096 Feb 8 20:52 /tmp
```

In a long listing, you can spot the **sticky** permissions by a lowercase **t** where you would normally expect the **x**(other execute permissions) to be. If the other does not have execute permissions, this will be replaced by an uppercase **T**.

Lastly, **setgid** on a directory means that files created in the directory will inherit the group affiliation from the directory, rather than inheriting it from the creating user. This is commonly used on group collaborative directories to automatically change a file from the default private group to the shared group.

In a long listing, you can spot the **setgid** permissions by a lowercase **s** where you would normally expect the **x**(group execute permissions) to be. If the group does not have execute permissions, this will be replaced by an uppercase **S**.

□

### Effects of special permissions on files and directories

| Special permission  | Effect on files                                                               | Effect on directories                                                                                                                   |
|---------------------|-------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| <b>u+s</b> (suid)   | File executes as the user that owns the file, not the user that ran the file. | No effect.                                                                                                                              |
| <b>g+s</b> (sgid)   | File executes as the group that owns the file.                                | Files newly created in the directory have their group owner set to match the group owner of the directory.                              |
| <b>o+t</b> (sticky) | No effect.                                                                    | Users with <b>write</b> on the directory can only remove files that they own; they cannot remove or force saves to files owned by other |

| Special permission | Effect on files | Effect on directories |
|--------------------|-----------------|-----------------------|
|                    |                 | users.                |

### Setting special permissions

- Symbolically: `setuid = u+s`; `setgid = g+s`; `sticky = o+t`
- Numerically (fourth preceding digit): `setuid = 4`; `setgid = 2`; `sticky = 1`

### Examples

- Add the setgid bit on **directory**:

```
[root@desktopX ~]# chmod g+s directory
```

- Set the setgid bit, and read/write/execute for user and group on **directory**:

```
[root@desktopX ~]# chmod 2770 directory
```

## Default file permissions

The default permissions for files are set by the processes that create them. For example, text editors create files so they are readable and writeable, but not executable, by everyone. The same goes for shell redirection. Additionally, binary executables are created executable by the compilers that create them. The **mkdir** command creates new directories with all permissions set—read, write, and execute.

Experience shows that these permissions are not typically set when new files and directories are created. This is because some of the permissions are cleared by the `umask` of the shell process.

The **umask** command without arguments will display the current value of the shell's `umask`:

```
[student@desktopX ~]$ umask
0002
```

Every process on the system has a `umask`, which is an octal bitmask that is used to clear the permissions of new files and directories that are created by the process. If a bit is set in the `umask`, then the corresponding permission is cleared in new files. For example, the previous `umask`, `0002`, clears the write bit for other users. The leading zeros indicate the special, user, and group permissions are not cleared. A `umask` of `077` clears all the group and other permissions of newly created files.

Use the **umask** command with a single numeric argument to change the `umask` of the current shell. The numeric argument should be an octal value corresponding to the new `umask` value. If it is less than 3 digits, leading zeros are assumed.

The system default `umask` values for Bash shell users are defined in the `/etc/profile` and `/etc/bashrc` files. Users can override the system defaults in their `.bash_profile` and `.bashrc` files.

In this example, please follow along with the next steps while your instructor demonstrates the effects of **umask** on new files and directories.

1. Create a new file and directory to see how the default `umask` affects permissions.

```
2. [student@desktopX ~]$ touch newfile1
3. [student@desktopX ~]$ ls -l newfile1
4. -rw-rw-r--. 1 student student 0 May 9 01:54 newfile1
5. [student@desktopX ~]$ mkdir newdir1
6. [student@desktopX ~]$ ls -ld newdir1
7. drwxrwxr-x. 2 student student 0 May 9 01:54 newdir1
```

8. Set the umask value to 0. This setting will not mask any of the permissions of new files. Create a new file and directory to see how this new umask affects permissions.

```
9. [student@desktopX ~]$ umask 0
10. [student@desktopX ~]$ touch newfile2
11. [student@desktopX ~]$ ls -l newfile2
12. -rw-rw-rw-. 1 student student 0 May 9 01:54 newfile2
13. [student@desktopX ~]$ mkdir newdir2
14. [student@desktopX ~]$ ls -ld newdir2
15. drwxrwxrwx. 2 student student 0 May 9 01:54 newdir2
```

16. Set the umask value to 007. This setting will mask all of the “other” permissions of new files.

```
17. [student@desktopX ~]$ umask 007
18. [student@desktopX ~]$ touch newfile3
19. [student@desktopX ~]$ ls -l newfile3
20. -rw-rw----. 1 student student 0 May 9 01:55 newfile3
21. [student@desktopX ~]$ mkdir newdir3
22. [student@desktopX ~]$ ls -ld newdir3
23. drwxrwx---. 2 student student 0 May 9 01:54 newdir3
```

24. Set the umask value to 027. This setting will mask write access for group members and all of the “other”permissions of new files.

```
25. [student@desktopX ~]$ umask 027
26. [student@desktopX ~]$ touch newfile4
27. [student@desktopX ~]$ ls -l newfile4
28. -rw-r-----. 1 student student 0 May 9 01:55 newfile4
29. [student@desktopX ~]$ mkdir newdir4
30. [student@desktopX ~]$ ls -ld newdir4
31. drwxr-x---. 2 student student 0 May 9 01:54 newdir4
```

32. Log in as **root** to change the default umask for unprivileged users to prohibit all access for users not in their group.

Modify **/etc/bashrc** and **/etc/profile** to change the default umask for Bash shell users. Since the default umask for unprivileged users is 0002, look for the **umask** command in these files that sets the umask to that value. Change them to set the umask to 007.

```
[root@desktopX ~]# less /etc/bashrc
```

```
You could check uidgid reservation validity in
/usr/share/doc/setup-*/uidgid file
if [$UID -gt 199] && ["`id -gn`" = "`id -un`"]; then
 umask 002
else
 umask 022
fi

Only display echos from profile.d scripts if we are no login shell
```

```
[root@desktopX ~]# vim /etc/bashrc
```

```
[root@desktopX ~]# less /etc/bashrc
```

```
You could check uidgid reservation validity in
/usr/share/doc/setup-*/uidgid file
if [$UID -gt 199] && ["`id -gn`" = "`id -un`"]; then
 umask 007
else
 umask 022
fi

Only display echos from profile.d scripts if we are no login shell
```

```
[root@desktopX ~]# less /etc/profile
```

```
You could check uidgid reservation validity in
/usr/share/doc/setup-*/uidgid file
if [$UID -gt 199] && ["`id -gn`" = "`id -un`"]; then
 umask 002
else
 umask 022
fi

for i in /etc/profile.d/*.sh ; do
```

```
[root@desktopX ~]# vim /etc/profile
```

```
[root@desktopX ~]# less /etc/profile

You could check uidgid reservation validity in
/usr/share/doc/setup-*/uidgid file

if [$UID -gt 199] && ["`id -gn`" = "`id -un`"]; then

 umask 007

else

 umask 022

fi

for i in /etc/profile.d/*.sh ; do
```

33. Log back in as **student** and confirm that the umask changes you made are persistent.

```
34. [student@desktopX ~]$ umask
35. 0007
```

#### Note

Other shells, such as **tcsh**, may have different system default initialization files in **/etc** and users' home directories.

#### References

**bash**(1), **ls**(1), **chmod**(1), and **umask**(1) man pages

## 6.6. Practice: Controlling New File Permissions and

### Ownership

#### Guided exercise

In this lab, you will control default permissions on new files using the **umask** command and **setgid** permission.

#### Outcomes

- Create a shared directory where new files are automatically owned by the group **atteam**.
- Experiment with various umask settings.
- Adjust default permissions for specific users.
- Confirm your adjustment is correct.

#### Before you begin...

Reset your serverX system. Run **lab permissions setup** to create the **alice** account. The password for **alice** is **password**.

1. Log in as **alice** on your **serverX** virtual machine and open a window with a Bash prompt. Use the **umask** command without arguments to display Alice's default umask value.

```
2. [alice@serverX ~]$ umask
3. 0002
```



4. Create a new directory `/tmp/shared` and a new file `/tmp/shared/defaults` to see how the default umask affects permissions.

```
5. [alice@serverX ~]$ mkdir /tmp/shared
6. [alice@serverX ~]$ ls -ld /tmp/shared
7. drwxrwxr-x. 2 alice alice 6 Jan 26 18:43 /tmp/shared
8. [alice@serverX ~]$ touch /tmp/shared/defaults
9. [alice@serverX ~]$ ls -l /tmp/shared/defaults
10. -rw-rw-r--. 1 alice alice 0 Jan 26 18:43 /tmp/shared/defaults
```

11. Change the group ownership of `/tmp/shared` to `ateam` and record the new ownership and permissions.

```
12. [alice@serverX ~]$ chown :ateam /tmp/shared
13. [alice@serverX ~]$ ls -ld /tmp/shared
14. drwxrwxr-x. 2 alice ateam 21 Jan 26 18:43 /tmp/shared
```

15. Create a new file in `/tmp/shared` and record the ownership and permissions.

```
16. [alice@serverX ~]$ touch /tmp/shared/alice3
17. [alice@serverX ~]$ ls -l /tmp/shared/alice3
18. -rw-rw-r--. 1 alice alice 0 Jan 26 18:46 /tmp/shared/alice3
```

19. Ensure the permissions of `/tmp/shared` cause files created in that directory to inherit the group ownership of `ateam`.

```
20. [alice@serverX ~]$ chmod g+s /tmp/shared
21. [alice@serverX ~]$ ls -ld /tmp/shared
22. drwxrwsr-x. 2 alice ateam 34 Jan 26 18:46 /tmp/shared
23. [alice@serverX ~]$ touch /tmp/shared/alice4
24. [alice@serverX ~]$ ls -l /tmp/shared/alice4
25. -rw-rw-r--. 1 alice ateam 0 Jan 26 18:48 /tmp/shared/alice4
```

26. Change the umask for `alice` such that new files are created with read-only access for the group and no access for other users. Create a new file and record the ownership and permissions.

```
27. [alice@serverX ~]$ umask 027
28. [alice@serverX ~]$ touch /tmp/shared/alice5
29. [alice@serverX ~]$ ls -l /tmp/shared/alice5
30. -rw-r-----. 1 alice ateam 0 Jan 26 18:48 /tmp/shared/alice5
```

31. Open a new Bash shell as `alice` and view the umask.

```
32. [alice@serverX ~]$ umask
33. 0002
```

34. Change the default umask for **alice** to prohibit all access for users not in their group.

```
35. [alice@serverX ~]# echo "umask 007" >> ~/.bashrc
36. [alice@serverX ~]# cat ~/.bashrc
37. # .bashrc
38.
39. # Source global definitions
40. if [-f /etc/bashrc]; then
41. . /etc/bashrc
42. fi
43.
44. # Uncomment the following line if you don't like systemctl's auto-paging feature:
45. # export SYSTEMD_PAGER=
46.
47. # User specific aliases and functions
48. umask 007
```

49. Log out and back into **serverX** as **alice** and confirm that the umask changes you made are persistent.

```
50. [alice@serverX ~]$ umask
```

```
0007
```

## 6.7. Lab: Controlling Access to Files with Linux File System

### Permissions

#### Performance checklist

In this lab, you will configure a system with directories for user collaboration.

#### Outcomes

- A directory on serverX called **/home/stooges** where these three users can work collaboratively on files.
- Only the user and group access, create, and delete files in **/home/stooges**. Files created in this directory should automatically be assigned a group ownership of **stooges**.
- New files created by users will not be accessible outside of the group.

#### Before you begin...

Reset your serverX system. Log into and set up your server system.

```
[student@serverX ~]$ lab permissions setup
```

Your serverX machine has three accounts, **curly**, **larry**, and **moe**, who are members of a group called **stooges**. The password for each account is **password**.

1. Open a terminal window and become root on serverX.

```
2. [student@serverX ~]$ su -
3. Password: redhat
```

```
[root@serverX ~]#
```

4. Create the **/home/stooges** directory.

```
[root@serverX ~]# mkdir /home/stooges
```

5. Change group permissions on the **/home/stooges** directory so it belongs to the **stooges** group.

```
[root@serverX ~]# chown :stooges /home/stooges
```

6. Set permissions on the **/home/stooges** directory so it is a set GID bit directory (2), the owner (7) and group (7) have full read/write/execute permissions, and other users have no permission (0) to the directory.

```
[root@serverX ~]# chmod 2770 /home/stooges
```

7. Check that the permissions were set properly.

```
8. [root@serverX ~]# ls -ld /home/stooges
```

```
drwxrws---. 2 root stooges 1024 Dec 9 1:38 /home/stooges
```

9. Modify the global login scripts so that normal users have a **umask** setting which prevents others from viewing or modifying new files and directories.

```
10. [root@serverX ~]# vim /etc/bashrc
11. [root@serverX ~]# vim /etc/profile
12. [root@serverX ~]# less /etc/bashrc
13. # You could check uidgid reservation validity in
14. # /usr/share/doc/setup-*/uidgid file
15. if [$UID -gt 199] && ["`id -gn`" = "`id -un`"]; then
16. umask 007
17. else
18. umask 022
19. fi
20.
```

```
21. for i in /etc/profile.d/*.sh ; do
```

22. When you finish, open a terminal window on serverX and run **lab permissions grade** to confirm you have done everything correctly.

```
[student@serverX ~]$ lab permissions grade
```

# Chapter 7. Monitoring and Managing Linux Processes

[7.1. Processes](#)

[7.2. Practice: Processes](#)

[7.3. Controlling Jobs](#)

[7.4. Practice: Background and Foreground Processes](#)

[7.5. Killing Processes](#)

[7.6. Practice: Killing Processes](#)

[7.7. Monitoring Process Activity](#)

[7.8. Practice: Monitoring Process Activity](#)

[7.9. Lab: Monitoring and Managing Linux Processes](#)

Introduction

| Overview          |                                                                                                                                                                                                                                                                                                                                             |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Goal</b>       | To evaluate and control processes running on a Red Hat Enterprise Linux system.                                                                                                                                                                                                                                                             |
| <b>Objectives</b> | <ul style="list-style-type: none"><li>○ List and interpret basic information about processes running on the system.</li><li>○ Control processes in the shell's session using bash job control.</li><li>○ Terminate and control processes using signals.</li><li>○ Monitor resource usage and system load due to process activity.</li></ul> |
| <b>Sections</b>   | <ul style="list-style-type: none"><li>○ Processes (and Practice)</li><li>○ Controlling Jobs (and Practice)</li></ul>                                                                                                                                                                                                                        |

## Overview

|            |                                                                                                                                         |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------|
|            | <ul style="list-style-type: none"><li>○ Killing Processes (and Practice)</li><li>○ Monitoring Process Activity (and Practice)</li></ul> |
| <b>Lab</b> | <ul style="list-style-type: none"><li>○ Monitoring and Managing Linux Processes</li></ul>                                               |

## 7.1. Processes

### Objectives

After completing this section, students should be able to:

- Define the life cycle of a process.
- Define process states.
- View and interpret process listings.

### What is a process?

#### Processes

A *process* is a running instance of a launched, executable program. A process consists of:

- an address space of allocated memory,
- security properties including ownership credentials and privileges,
- one or more execution threads of program code, and
- the process state.

The *environment* of a process includes:

- local and global variables,
- a current scheduling context, and
- allocated system resources, such as file descriptors and network ports.

An existing (*parent*) process duplicates its own address space (**fork**) to create a new (*child*) process structure. Every new process is assigned a unique *process ID* (PID) for tracking and security. The PID and the *parent's process ID* (PPID) are elements of the new process environment. Any process may create a child process. All processes are descendants of the first system process, which is **systemd**(1) on a Red Hat Enterprise Linux 7 system.

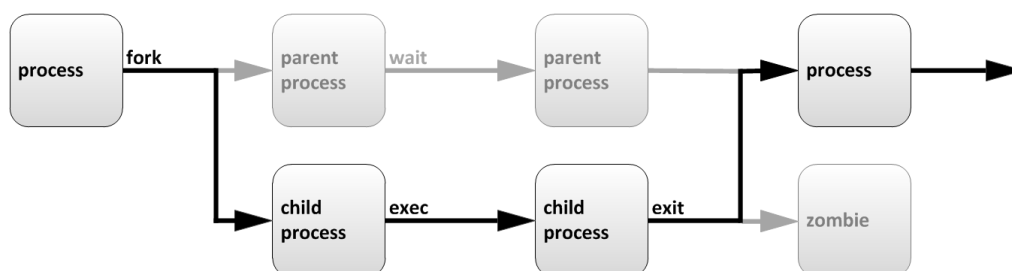


Figure 7.1: Process life cycle

Through the **fork** routine, a child process inherits security identities, previous and current file descriptors, port and resource privileges, environment variables, and program code. A child process may then **exec** its own program code. Normally, a parent process *sleeps* while the child process runs, setting a request (**wait**) to be signaled when the child completes. Upon exit, the child process has already closed or discarded its resources and environment; the remainder is referred to as a *zombie*. The parent, signaled awake when the child exited, cleans the remaining structure, then continues with its own program code execution.

## Process states

In a multitasking operating system, each CPU (or CPU core) can be working on one process at a single point in time. As a process runs, its immediate requirements for CPU time and resource allocation change. Processes are assigned a *state*, which changes as circumstances require.

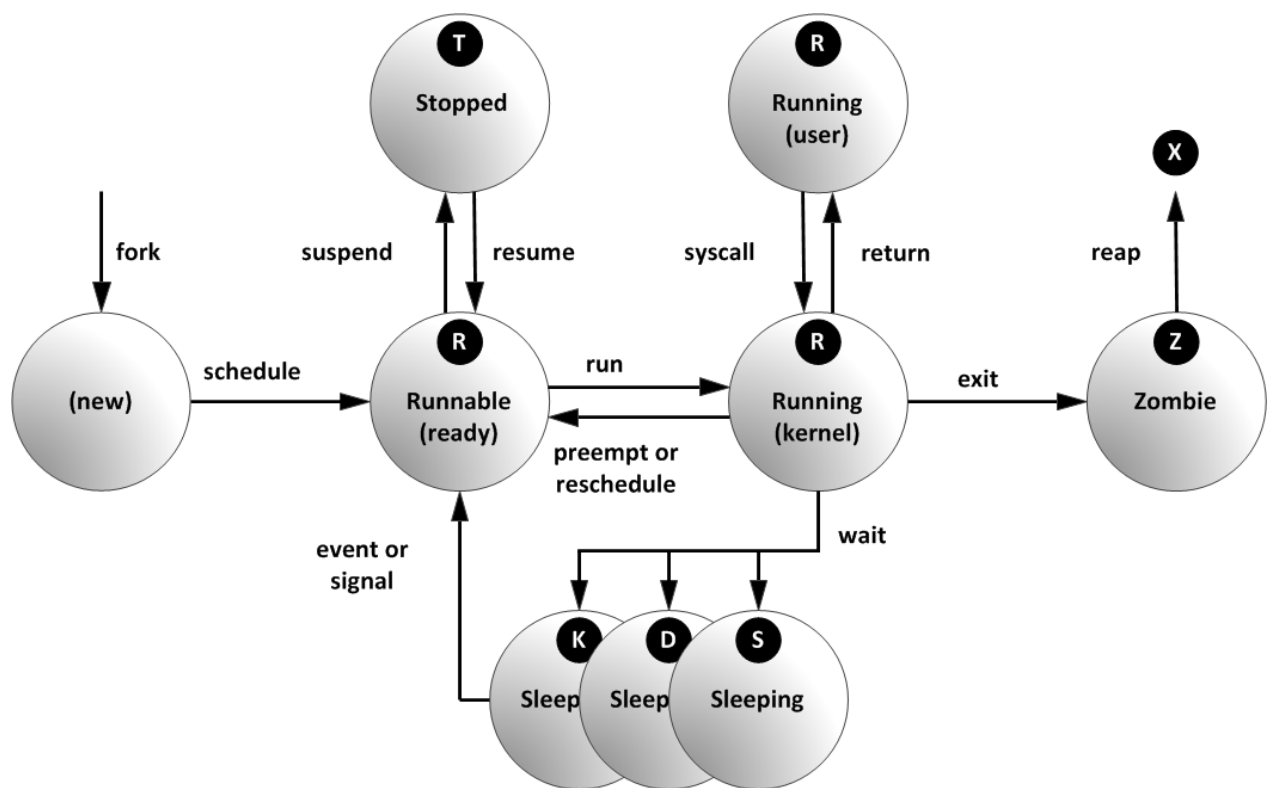


Figure 7.2: Linux process states

The Linux process states are illustrated in the previous diagram and described in the following table.

□

### Linux process states

| Name    | Flag     | Kernel-defined state name and description                                                                                                                                                                                          |
|---------|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Running | <b>R</b> | TASK_RUNNING: The process is either executing on a CPU or waiting to run. Process can be executing user routines or kernel routines (system calls), or be queued and ready when in the <i>Running</i> (or <i>Runnable</i> ) state. |

| Name     | Flag     | Kernel-defined state name and description                                                                                                                                                                                                 |
|----------|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Sleeping | <b>S</b> | TASK_INTERRUPTIBLE: The process is waiting for some condition: a hardware request, system resource access, or signal. When an event or signal satisfies the condition, the process returns to <i>Running</i> .                            |
|          | <b>D</b> | TASK_UNINTERRUPTIBLE: This process is also <i>Sleeping</i> , but unlike <b>S</b> state, will not respond to delivered signals. Used only under specific conditions in which process interruption may cause an unpredictable device state. |
|          | <b>K</b> | TASK_KILLABLE: Identical to the uninterruptible <b>D</b> state, but modified to allow the waiting task to respond to a signal to be killed (exited completely). Utilities frequently display <i>Killable</i> processes as <b>D</b> state. |
| Stopped  | <b>T</b> | TASK_STOPPED: The process has been <i>Stopped</i> (suspended), usually by being signaled by a user or another process. The process can be continued (resumed) by another signal to return to <i>Running</i> .                             |
|          | <b>T</b> | TASK_TRACED: A process that is being debugged is also temporarily <i>Stopped</i> and shares the same <b>T</b> state flag.                                                                                                                 |
| Zombie   | <b>Z</b> | EXIT_ZOMBIE: A child process signals its parent as it exits. All resources except for the process identity (PID) are released.                                                                                                            |
|          | <b>X</b> | EXIT_DEAD: When the parent cleans up ( <i>reaps</i> ) the remaining child process structure, the process is now released completely. This state will never be observed in process-listing utilities.                                      |

## Listing processes

The **ps** command is used for listing current processes. The command can provide detailed process information, including:

- the user identification (UID) which determines process privileges,
- the unique process identification (PID),
- the CPU and real time already expended,
- how much memory the process has allocated in various locations,
- the location of process **STDOUT**, known as the *controlling terminal*, and
- the current process state.

### Important

The Linux version of **ps** supports three option formats, including:

- UNIX (POSIX) options, which may be grouped and must be preceded by a dash,

- BSD options, which may be grouped and must not be used with a dash, and
- GNU long options, which are preceded by two dashes.

For example, **ps -aux** is not the same as **ps aux**.

A common display listing (options **aux**) displays all processes, with columns in which users will be interested, and includes processes without a controlling terminal. A long listing (options **lax**) provides more technical detail, but may display faster by avoiding the username lookup. The similar UNIX syntax uses the options **-ef** to display all processes.

```
[student@serverX ~]$ ps aux
```

| USER | PID | %CPU | %MEM | VSZ   | RSS  | TTY | STAT | START | TIME | COMMAND               |
|------|-----|------|------|-------|------|-----|------|-------|------|-----------------------|
| root | 1   | 0.1  | 0.1  | 51648 | 7504 | ?   | Ss   | 17:45 | 0:03 | /usr/lib/systemd/syst |
| root | 2   | 0.0  | 0.0  | 0     | 0    | ?   | S    | 17:45 | 0:00 | [kthreadd]            |
| root | 3   | 0.0  | 0.0  | 0     | 0    | ?   | S    | 17:45 | 0:00 | [ksoftirqd/0]         |
| root | 5   | 0.0  | 0.0  | 0     | 0    | ?   | S<   | 17:45 | 0:00 | [kworker/0:0H]        |
| root | 7   | 0.0  | 0.0  | 0     | 0    | ?   | S    | 17:45 | 0:00 | [migration/0]         |

-- output truncated --

```
[student@serverX ~]$ ps lax
```

| F | UID | PID | PPID | PRI  | NI  | VSZ   | RSS  | WCHAN  | STAT | TTY | TIME | COMMAND           |
|---|-----|-----|------|------|-----|-------|------|--------|------|-----|------|-------------------|
| 4 | 0   | 1   | 0    | 20   | 0   | 51648 | 7504 | ep_pol | Ss   | ?   | 0:03 | /usr/lib/systemd/ |
| 1 | 0   | 2   | 0    | 20   | 0   | 0     | 0    | kthrea | S    | ?   | 0:00 | [kthreadd]        |
| 1 | 0   | 3   | 2    | 20   | 0   | 0     | 0    | smpboo | S    | ?   | 0:00 | [ksoftirqd/0]     |
| 1 | 0   | 5   | 2    | 0    | -20 | 0     | 0    | worker | S<   | ?   | 0:00 | [kworker/0:0H]    |
| 1 | 0   | 7   | 2    | -100 | -   | 0     | 0    | smpboo | S    | ?   | 0:00 | [migration/0]     |

-- output truncated --

```
[student@serverX ~]$ ps -ef
```

| UID  | PID | PPID | C | STIME | TTY | TIME     | CMD                                    |
|------|-----|------|---|-------|-----|----------|----------------------------------------|
| root | 1   | 0    | 0 | 17:45 | ?   | 00:00:03 | /usr/lib/systemd/systemd --switched-ro |



```

root 2 0 0 17:45 ? 00:00:00 [kthreadd]

root 3 2 0 17:45 ? 00:00:00 [ksoftirqd/0]

root 5 2 0 17:45 ? 00:00:00 [kworker/0:0H]

root 7 2 0 17:45 ? 00:00:00 [migration/0]

-- output truncated --

```

By default, **ps** with no options selects all processes with the same *effective user ID* (EUID) as the current user and associated with the same terminal where **ps** was invoked.

- Processes in brackets (usually at the top) are scheduled kernel threads.
- Zombies show up in a **ps** listing as *exiting* or *defunct*.
- **ps** displays once. Use **top(1)** for a repetitive update process display.
- **ps** can display in tree format to view parent/child relationships.
- The default output is unsorted. Display order matches that of the system process table, which reuses table rows as processes die and new ones are created. Output may appear chronological, but is not guaranteed unless explicit **-O** or **--sort** options are used.

**References**

**info libc signal** (*GNU C Library Reference Manual*)

- Section 24: Signal Handling

**info libc processes** (*GNU C Library Reference Manual*)

- Section 26: Processes

**ps(1)** and **signal(7)** man pages

## 7.2. Practice: Processes

### Quiz

Match the following items to their counterparts in the table.

| Description                                                 | State |
|-------------------------------------------------------------|-------|
| Process has been stopped temporarily.                       | T     |
| Process has just been terminated.                           | Z     |
| Process is scheduled but is not yet on a CPU.               | R     |
| Process is waiting for I/O.                                 | S     |
| Process is uninterruptibly waiting for a device to respond. | D     |
| Process is at a prompt, needing user input.                 | S     |
| Process is executing a system call.                         | R     |

## 7.3. Controlling Jobs

### Objectives

After completing this section, students should be able to:

- Explain the terms foreground, background, and controlling terminal.
- Use job control to manage multiple command-line tasks.

### Jobs and sessions

#### Controlling jobs

*Job control* is a command shell feature allowing a single shell instance to run and manage multiple commands. Without job control, a parent shell forks a child process to run a command, sleeping until the child process exits. When the shell prompt redisplay, the parent shell has returned. With job control, commands can be selectively suspended, resumed, and run asynchronously, allowing the shell to return for additional commands while child processes run.

A *foreground* process is a command running in a terminal window. The terminal's device ID (**tty**) is the process's *controlling terminal*. Foreground processes receive keyboard-generated input and signals and are allowed to read from or write to the terminal (e.g., via **stdin** and **stdout**).

A process *session* is created when a terminal or console first opens (e.g., at login or by invoking a new **Terminal** instance). All processes (e.g., the first command shell, its children, and pipelines) initiated from that terminal share the same *session ID*. Within a session, only one process can be in the foreground at a time.

A *background* process is started *without* a controlling terminal because it has no need for terminal interaction. In **ps** listing, such processes (e.g., service daemons and kernel process threads) display a question mark (?) in the **TTY** column. Background processes which (improperly) attempt to read from or write to the terminal may be suspended.

### Running jobs in the background

Any command can be started in the background by appending an ampersand (**&**) to the command line. The **bash** shell displays a *job number* (unique to the session) and the PID of the new child process. The command shell does not wait for the child process and redisplay the shell prompt.

```
[student@serverX ~]$ sleep 10000 &
[1] 5947
[student@serverX ~]$
```

#### Note

An ampersand will background only the last command in a pipeline, unless the command set is

surrounded with parentheses. The PID returned is from the pipeline's last process. All processes in the pipeline are now members of the same job.

```
[student@serverX ~]$ (example_command | sort | mail -s "Sort
output") &

[1] 5998
```

The **bash** command shell tracks jobs, per session, in a table displayed with the **jobs** command.

```
[student@serverX ~]$ jobs

[1]+ Running sleep 10000 &

[student@serverX ~]$
```

Background jobs can reconnect to the controlling terminal by being brought to the foreground using the **fg** command with the job ID (*%job number*).

```
[student@serverX ~]$ fg %1

sleep 10000

_
```

The example **sleep** command is now running on the controlling terminal. The command shell is again asleep, waiting for this child process to exit. To resend to the background, or to send any command in which the trailing ampersand was not originally included, send the keyboard-generated *suspend* request (**Ctrl-z**) to the process.

```
sleep 10000

^Z

[1]+ Stopped sleep 10000

[student@serverX ~]$
```

The **suspend** takes effect immediately. The job is placed in the background. Pending output and keyboard typeahead are discarded.

The **ps** option **j** displays job information, including the initial command shell of each session. Since the **sleep** example command is currently suspended, the state flag displayed is **T**.

```
[student@serverX ~]$ ps j

PPID PID PGID SID TTY TPGID STAT UID TIME COMMAND
2764 2768 2768 2768 pts/0 6377 Ss 1000 0:00 /bin/bash
2768 5947 5947 2768 pts/0 6377 T 1000 0:00 sleep 10000
2768 6377 6377 2768 pts/0 6377 R+ 1000 0:00 ps j

[student@serverX ~]$
```

To restart the process in the background, use the **bg** command with the same job ID.

```
[student@serverX ~]$ bg %1
[1]+ sleep 10000 &
[student@serverX ~]$
```

The command shell will warn a user who attempts to exit a terminal window (session) with suspended jobs. If the user tries exiting again immediately, the suspended jobs are killed.

#### References

Additional information may be available in the chapter on viewing system processes in the *Red Hat Enterprise Linux System Administrator's Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

**libc** info page (*GNU C Library Reference Manual*)

- Section 24: Signal Handling
- Section 26: Processes

**bash(1)**, **builtins(1)**, **ps(1)**, **sleep(1)** man pages

## 7.4. Practice: Background and Foreground Processes

### Guided exercise

In this lab, students will start, suspend, and reconnect to multiple processes using job control.

#### Outcomes:

Practice suspending and restarting user processes.

#### Before you begin...

Log in as student to serverX. Begin in student's home directory.

1. Open two terminal windows, side by side, to be referred to as *left* and *right*.
2. In the left window, start a process that continuously appends the word "rock" and a space to the file **~/outfile** at one-second intervals. The complete command set must be contained in parentheses for job control to interpret the set as a single job.

```
[student@serverX ~]$ (while true; do echo -n "rock " >>
~/outfile; sleep 1; done)
```

3. In the right window, use **tail** to confirm that the new process is writing to the file.

```
[student@serverX ~]$ tail -f ~/outfile
```

4. In the left window, suspend the running process. The shell returns the job ID in square brackets. In the right window, confirm that the process output has stopped.

```
[student@serverX ~]$ Ctrl-z
```

5. In the left window, view the **jobs** list. The **+** denotes the *current job*. Restart the job in the background. In the right window, confirm that the process output is again active.

```
6. [student@serverX ~]$ jobs
7. [1]+ Stopped (while true; do
8. echo -n "rock " >> ~/outfile; sleep 1;
9. done)
10. [student@serverX ~]$ bg
```

```
[student@serverX ~]$ jobs
```

11. In the left window, start two more processes to append to the same output file. Replace "rock" with "paper," and then with "scissors." To properly background the process, the complete command set must be contained in parentheses and ended with an ampersand.

```
12. [student@serverX ~]$ (while true; do echo -n "paper " >>
~/outfile; sleep 1; done) &
```

```
[student@serverX ~]$ (while true; do echo -n "scissors
" >> ~/outfile; sleep 1; done) &
```

13. In the left window, view **jobs** to see all three processes "Running". In the right window, confirm that all three processes are appending to the file.

```
[student@serverX ~]$ jobs
```

14. Using only commands previously learned, suspend the "rock" process. In the left window, foreground the job, using the job ID determined from the **jobs** listing, then suspend it using Ctrl-z. Confirm that the "rock" process is "Stopped". In the right window, confirm that "rock" output is no longer active.

```
15. [student@serverX ~]$ jobs
16. [student@serverX ~]$ fg %number
```

```
[student@serverX ~]$ Ctrl-z
```

17. End the "paper" process. In the left window, foreground the job, then terminate it using Ctrl-c. Confirm that the "paper" process has disappeared. In the right window, confirm that "paper" output is no longer active.

```
18. [student@serverX ~]$ jobs
19. [student@serverX ~]$ fg %number
```

```
[student@serverX ~]$ Ctrl-c
```

20. In the left window, view the remaining jobs using **ps**. The suspended job has state **T**. The other background job is sleeping (**S**), since **ps** is "on cpu" (**R**) while displaying.

```
21. [student@serverX ~]$ ps j
22. PPID PID PGID SID TTY TPGID STAT UID TIME COMMAND
23. 4489 6223 6223 6223 pts/1 12918 Ss 1000 0:00 bash
24. 4489 6237 6237 6237 pts/2 9782 Ss 1000 0:00 bash
25. 6237 9782 9782 6237 pts/2 9782 S+ 1000 0:00 tail -f /home/student/o
26. 7360 9856 7360 6223 pts/1 12918 T 1000 0:00 sleep 1
27. 7395 12916 7395 6223 pts/1 12918 S 1000 0:00 sleep 1
```

```
6223 12918 12918 6223 pts/1 12918 R+ 1000 0:00 ps j
```

28. Stop the remaining two jobs. In the left window, foreground either job. Terminate it using **Ctrl-c**. Repeat with the remaining job. The "Stopped" job temporarily restarts when foregrounded. Confirm that no jobs remain and that output has stopped.

```
29. [student@serverX ~]$ fg %number
30. [student@serverX ~]$ Ctrl-c
31. [student@serverX ~]$ fg %number
32. [student@serverX ~]$ Ctrl-c
```

```
[student@serverX ~]$ jobs
```

33. In the right window, stop the **tail** command. Close extra terminal windows.

```
[student@serverX ~]$ Ctrl-c
```

## 7.5. Killing Processes

### Objectives

After completing this section, students should be able to:

- Use commands to kill and communicate with processes.
- Define the characteristics of a daemon process.
- End user sessions and processes.

### Process control using signals

#### *Killing processes*

A signal is a software interrupt delivered to a process. Signals report events to an executing program. Events that generate a signal can be an error, external event (e.g., I/O request or expired timer), or

by explicit request (e.g., use of a signal-sending command or by keyboard sequence).

The following table lists the fundamental signals used by system administrators for routine process management. Refer to signals by either their short (**HUP**) or proper (**SIGHUP**) name.

□

#### Fundamental process management signals

| Signal number | Short name | Definition         | Purpose                                                                                                                                                        |
|---------------|------------|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1             | HUP        | Hangup             | Used to report termination of the controlling process of a terminal. Also used to request process reinitialization (configuration reload) without termination. |
| 2             | INT        | Keyboard interrupt | Causes program termination. Can be blocked or handled. Sent by typing <b>INTR</b> character ( <b>Ctrl-c</b> ).                                                 |
| 3             | QUIT       | Keyboard quit      | Similar to <b>SIGINT</b> , but also produces a process dump at termination. Sent by typing <b>QUIT</b> character ( <b>Ctrl-\</b> ).                            |
| 9             | KILL       | Kill, unblockable  | Causes abrupt program termination. Cannot be blocked, ignored, or handled; always fatal.                                                                       |
| 15<br>default | TERM       | Terminate          | Causes program termination. Unlike <b>SIGKILL</b> , can be blocked, ignored, or handled. The polite way to ask a program to terminate; allows self-cleanup.    |
| 18            | CONT       | Continue           | Sent to a process to resume if stopped. Cannot be blocked. Even if handled, always resumes the process.                                                        |
| 19            | STOP       | Stop, unblockable  | Suspends the process. Cannot be blocked or handled.                                                                                                            |
| 20            | TSTP       | Keyboard stop      | Unlike <b>SIGSTOP</b> , can be blocked, ignored, or handled. Sent by typing <b>SUSP</b> character ( <b>Ctrl-z</b> ).                                           |

#### Note

Signal numbers vary on different Linux hardware platforms, but signal names and meanings are standardized. For command use, it is advised to use signal names instead of numbers. The numbers discussed in this section are for Intel x86 systems.

Each signal has a default action, usually one of the following:

Term — Cause a program to terminate (exit) at once.

Core — Cause a program to save a memory image (core dump), then terminate.

Stop — Cause a program to stop executing (suspend) and wait to continue (resume).

Programs can be prepared for expected event signals by implementing handler routines to ignore, replace, or extend a signal's default action.

### Commands for sending signals by explicit request

Users signal their current foreground process by typing a keyboard control sequence to suspend (**Ctrl-z**), kill (**Ctrl-c**), or core dump (**Ctrl-\**) the process. To signal a background process or processes in a different session requires a signal-sending command.

Signals can be specified either by name (e.g., **-HUP** or **-SIGHUP**) or by number (e.g., **-1**). Users may kill their own processes, but root privilege is required to kill processes owned by others.

- The **kill** command sends a signal to a process by ID. Despite its name, the kill command can be used for sending any signal, not just those for terminating programs.

```
[student@serverX ~]$ kill PID
[student@serverX ~]$ kill -signal PID
[student@serverX ~]$ kill -1
1) SIGHUP 2) SIGINT 3) SIGQUIT 4) SIGILL 5) SIGTRAP
6) SIGABRT 7) SIGBUS 8) SIGFPE 9) SIGKILL 10) SIGUSR1
11) SIGSEGV 12) SIGUSR2 13) SIGPIPE 14) SIGALRM 15) SIGTERM
16) SIGSTKFLT 17) SIGCHLD 18) SIGCONT 19) SIGSTOP 20) SIGTSTP
```

```
-- output truncated --
```

- Use **killall** to send a signal to one or more processes matching selection criteria, such as a command name, processes owned by a specific user, or all system-wide processes.

```
[student@serverX ~]$ killall command_pattern
[student@serverX ~]$ killall -signal command_pattern
```

```
[root@serverX ~]# killall -signal -u username command_pattern
```

- The **pkill** command, like **killall**, can signal multiple processes. **pkill** uses advanced selection criteria, which can include combinations of:

Command — Processes with a pattern-matched command name.

UID — Processes owned by a Linux user account, effective or real.

GID — Processes owned by a Linux group account, effective or real.

Parent — Child processes of a specific parent process.

Terminal — Processes running on a specific controlling terminal.



- `[student@serverX ~]$ pkill command_pattern`
- `[student@serverX ~]$ pkill -signal command_pattern`
- `[root@serverX ~]# pkill -G GID command_pattern`
- `[root@serverX ~]# pkill -P PPID command_pattern`
- `[root@serverX ~]# pkill -t terminal_name -U UID command_pattern`

## Logging users out administratively

The `w` command views users currently logged into the system and their cumulative activities. Use the `TTY` and `FROM` columns to determine the user's location.

All users have a controlling terminal, listed as `pts/N` while working in a graphical environment window (pseudo-terminal) or `ttyN` on a system console, alternate console, or other directly connected terminal device. Remote users display their connecting system name in the `FROM` column when using the `-f` option.

```
[student@serverX ~]$ w -f
12:43:06 up 27 min, 5 users, load average: 0.03, 0.17, 0.66
USER TTY FROM LOGIN@ IDLE JCPU PCPU WHAT
student :0 :0 12:20 ?xdm? 1:10 0.16s gdm-session-wor
student pts/0 :0 12:20 2.00s 0.08s 0.01s w -f
root tty2 :0 12:26 14:58 0.04s 0.04s -bash
bob tty3 :0 12:28 14:42 0.02s 0.02s -bash
student pts/1 desktop2.example.12:41 1:07 0.03s 0.03s -bash
[student@serverX ~]$
```

Discover how long a user has been on the system by viewing the session login time. For each session, CPU resources consumed by current jobs, included background tasks and children, are in the `JCPU` column. Current foreground process CPU consumption is in the `PCPU` column.

Users may be forced off a system for security violations, resource overallocation, or other administrative need. Users are expected to quit unnecessary applications, close unused command shells, and exit login sessions when requested.

When situations occur in which users cannot be contacted or have unresponsive sessions, runaway resource consumption, or improper system access, their sessions may need to be administratively terminated using signals.

### Important

Although `SIGTERM` is the default signal, `SIGKILL` is a commonly misused administrator favorite. Since the `SIGKILL` signal cannot be handled or ignored, it is always fatal. However, it forces termination without allowing the killed process to run self-cleanup routines. It is recommended to

send **SIGTERM** first, then retry with **SIGKILL** only if a process fails to respond.

Processes and sessions can be individually or collectively signaled. To terminate all processes for one user, use the **pkill** command. Because the initial process in a login session (session leader) is designed to handle session termination requests and ignore unintended keyboard signals, killing all of a user's processes and login shells requires using the **SIGKILL** signal.

```
[root@serverX ~]#_pgrep -l -u bob
6964 bash
6998 sleep
6999 sleep

7000 sleep

[root@serverX ~]#_pkill -SIGKILL -u bob

[root@serverX ~]#_pgrep -l -u bob

[root@serverX ~]#_
```

When processes requiring attention are in the same login session, it may not be necessary to kill all of a user's processes. Determine the controlling terminal for the session using the **w** command, then kill only processes which reference the same terminal ID. Unless **SIGKILL** is specified, the session leader (here, the **bash** login shell) successfully handles and survives the termination request, but all other session processes are terminated.

```
[root@serverX ~]#_pgrep -l -u bob
7391 bash
7426 sleep
7427 sleep

7428 sleep

[root@serverX ~]#_w -h -u bob
bob tty3 18:37 5:04 0.03s 0.03s -bash

[root@serverX ~]#_pkill -t tty3

[root@serverX ~]#_pgrep -l -u bob

7391 bash

[root@serverX ~]#_pkill -SIGKILL -t tty3
```

```
[root@serverX ~]# _pgrep -l -u bob
```

```
[root@serverX ~]# _
```

The same selective process termination can be applied using parent and child process relationships. Use the `pstree` command to view a process tree for the system or a single user. Use the parent process's PID to kill all children they have created. This time, the parent `bash` login shell survives because the signal is directed only at its child processes.

```
[root@serverX ~]# pstree -p bob
```

```
bash(8391) └─sleep(8425)
 └─sleep(8426)
 └─sleep(8427)
```

```
[root@serverX ~]# pkill -P 8391
```

```
[root@serverX ~]# pgrep -l -u bob
```

```
bash(8391)
```

```
[root@serverX ~]# pkill -SIGKILL -P 8391
```

```
[root@serverX ~]# pgrep -l -u bob
```

```
bash(8391)
```

```
[root@serverX ~]#
```

## References

**info libc signal** (GNU C Library Reference Manual)

- Section 24: Signal Handling

**info libc processes** (GNU C Library Reference Manual)

- Section 26: Processes

**kill(1), killall(1), pgrep(1), pkill(1), pstree(1), signal(7), and w(1) man pages**

## 7.6. Practice: Killing Processes

### Guided exercise

In this lab, students will use keyboard sequences and signals to manage and stop processes.

#### Outcomes:

Experience with observing the results of starting and stopping multiple shell processes.

Before you begin...

Log in as student to serverX. Start in your home directory.

- Open two terminal windows, side by side, to be referred to as *left* and *right*.
- In the left window, start three processes that append text to an output file at one-second

intervals. To properly background each process, the complete command set must be contained in parentheses and ended with an ampersand.

- **[student@serverX ~]\$ (while true; do echo -n "game " >> ~/outfile; sleep 1; done) &**
- **[student@serverX ~]\$ (while true; do echo -n "set " >> ~/outfile; sleep 1; done) &**

```
[student@serverX ~]$ (while true; do echo -n "match " >>
~/outfile; sleep 1; done) &
```

- In the right window, use **tail** to confirm that all three processes are appending to the file. In the left window, view **jobs** to see all three processes "Running".

- **[student@serverX ~]\$ tail -f ~/outfile**
- **[student@serverX ~]\$ jobs**
- [1] Running ( while true; do
- echo -n "game " >> ~/outfile; sleep 1;
- done ) &
- [2]- Running ( while true; do
- echo -n "set " >> ~/outfile; sleep 1;
- done ) &
- [3]+ Running ( while true; do
- echo -n "match " >> ~/outfile; sleep 1;
- done ) &

- Suspend the "game" process using signals. Confirm that the "game" process is "Stopped". In the right window, confirm that "game" output is no longer active.

- **[student@serverX ~]\$ kill -SIGSTOP %number**

```
[student@serverX ~]$ jobs
```

- Terminate the "set" process using signals. Confirm that the "set" process has disappeared. In the right window, confirm that "set" output is no longer active.

- **[student@serverX ~]\$ kill -SIGTERM %number**

```
[student@serverX ~]$ jobs
```

- Resume the "game" process using signals. Confirm that the "game" process is "Running". In the right window, confirm that "game" output is again active.

- **[student@serverX ~]\$ kill -SIGCONT %number**

```
[student@serverX ~]$ jobs
```

- Terminate the remaining two jobs. Confirm that no jobs remain and that output has stopped. From the left window, terminate the right window's `tail` command.

Close extra terminal windows.

```
[student@serverX ~]$ kill -SIGTERM %number
[student@serverX ~]$ kill -SIGTERM %number
[student@serverX ~]$ jobs
[student@serverX ~]$ pkill -SIGTERM tail
[student@serverX ~]$
```

## 7.7. Monitoring Process Activity

### Objectives

After completing this section, students should be able to:

- Interpret uptime and load averages.
- Monitor real-time processes.

### Load average

#### Monitoring process activity

The Linux kernel calculates a *load average* metric as an *exponential moving average* of the *load number*, a cumulative CPU count of active system resource requests.

- *Active requests* are counted from per-CPU queues for running threads and threads waiting for I/O, as the kernel tracks process resource activity and corresponding process state changes.
- *Load number* is a calculation routine run every five seconds by default, which accumulates and averages the active requests into a single number for all CPUs.
- *Exponential moving average* is a mathematical formula to smooth out trending data highs and lows, increase current activity significance, and decrease aging data quality.
- *Load average* is the load number calculation routine result. Collectively, it refers to the three displayed values of system activity data averaged for the last 1, 5, and 15 minutes.

#### Understanding the Linux load average calculation

The load average represents the perceived system load over a time period. Linux implements the load average calculation as a representation of expected service wait times, not only for CPU but also for disk and network I/O.

- Linux counts not only processes, but threads individually, as separate tasks. CPU request queues for running threads (*nr\_running*) and threads waiting for I/O resources (*nr\_iowait*) reasonably correspond to process states **R**(*Running*) and **D** (*Uninterruptable Sleeping*). Waiting for I/O includes tasks sleeping for expected disk and network responses.

- The load number is a global counter calculation, which is sum-totaled for all CPUs. Since tasks returning from sleep may reschedule to different CPUs, accurate per-CPU counts are difficult, but an accurate cumulative count is assured. Displayed load averages represent all CPUs.
- Linux counts each physical CPU core and microprocessor hyperthread as separate execution units, logically represented and referred to as individual CPUs. Each CPU has independent request queues. View `/proc/cpuinfo` for the kernel representation of system CPUs.

```

[student@serverX ~]$ grep "model name" /proc/cpuinfo
model name : Intel(R) Core(TM) i5 CPU M 520 @ 2.40GHz
model name : Intel(R) Core(TM) i5 CPU M 520 @ 2.40GHz
model name : Intel(R) Core(TM) i5 CPU M 520 @ 2.40GHz
model name : Intel(R) Core(TM) i5 CPU M 520 @ 2.40GHz
[student@serverX ~]$ grep "model name" /proc/cpuinfo
| wc -l

```

```
4
```

- Some UNIX systems only considered CPU utilization or run queue length to indicate system load. Since a system with idle CPUs can experience extensive waiting due to busy disk or network resources, I/O consideration is included in the Linux load average. When experiencing high load averages with minimal CPU activity, examine the disk and network activity.

#### Interpreting displayed load average values

The three values represent the weighted values over the last 1, 5, and 15 minutes. A quick glance can indicate whether system load appears to be increasing or decreasing. Calculate the approximate *per-CPU* load value to determine whether the system is experiencing significant waiting.

- `top`, `uptime`, `w`, and `gnome-system-monitor` display load average values.

```
[student@serverX ~]$ uptime
```

```
15:29:03 up 14 min, 2 users, load average: 2.92, 4.48, 5.20
```

- Divide the displayed load average values by the number of logical CPUs in the system. A value below 1 indicates satisfactory resource utilization and minimal wait times. A value above 1 indicates resource saturation and some amount of service waiting times.

```

From /proc/cpuinfo, system has four logical CPUs, so divide by 4:
#
load average: 2.92, 4.48, 5.20
#
divide by number of logical CPUs: 4 4 4
#

#
per-CPU load average: 0.73 1.12 1.30
#
This system's load average appears to be decreasing.

```

- # With a load average of 2.92 on four CPUs, all CPUs were in use ~73% of the time.
- # During the last 5 minutes, the system was overloaded by ~12%.

```
During the last 15 minutes, the system was overloaded by ~30%.
```

- An idle CPU queue has a load number of 0. Each ready and waiting thread adds a count of 1. With a total queue count of 1, the resource (CPU, disk, or network) is in use, but no requests spend time waiting. Additional requests increment the count, but since many requests can be processed within the time period, resource *utilization* increases, but not *wait times*.
- Processes sleeping for I/O due to a busy disk or network resource are included in the count and increase the load average. While not an indication of CPU utilization, the queue count still indicates that users and programs are waiting for resource services.
- Until resource saturation, a load average will remain below 1, since tasks will seldom be found waiting in queue. Load average only increases when resource saturation causes requests to remain queued and counted by the load calculation routine. When resource utilization approaches 100%, each additional request starts experiencing service wait time.

## Real-time process monitoring

The **top** program is a dynamic view of the system's processes, displaying a summary header followed by a process or thread list similar to **ps** information. Unlike the static **ps** output, **top** continuously refreshes at a configurable interval, and provides capabilities for column reordering, sorting, and highlighting. User configurations can be saved and made persistent.

Default output columns are recognizable from other resource tools:

- The process ID (**PID**).
- User name (**USER**) is the process owner.
- Virtual memory (**VIRT**) is all memory the process is using, including the resident set, shared libraries, and any mapped or swapped memory pages. (Labeled **VSZ** in the **ps** command.)
- Resident memory (**RES**) is the physical memory used by the process, including any resident shared objects. (Labeled **RSS** in the **ps** command.)
- Process state (**S**) displays as:
  - **D** = Uninterruptable Sleeping
  - **R** = Running or Runnable
  - **S** = Sleeping
  - **T** = Stopped or Traced
  - **Z** = Zombie
- CPU time (**TIME**) is the total processing time since the process started. May be toggled to include cumulative time of all previous children.
- The process command name (**COMMAND**).

□

## Fundamental keystrokes in top

| Key              | Purpose                                                                                |
|------------------|----------------------------------------------------------------------------------------|
| ? or h           | Help for interactive keystrokes.                                                       |
| l, t, m          | Toggles for load, threads, and memory header lines.                                    |
| 1                | Toggle showing individual CPUs or a summary for all CPUs in header.                    |
| s <sup>(1)</sup> | Change the refresh (screen) rate, in decimal seconds (e.g., 0.5, 1, 5).                |
| b                | Toggle reverse highlighting for <i>Running</i> processes; default is bold only.        |
| B                | Enables use of bold in display, in the header, and for <i>Running</i> processes.       |
| H                | Toggle threads; show process summary or individual threads.                            |
| u, U             | Filter for any user name (effective, real).                                            |
| M                | Sorts process listing by memory usage, in descending order.                            |
| P                | Sorts process listing by processor utilization, in descending order.                   |
| k <sup>(1)</sup> | Kill a process. When prompted, enter <b>PID</b> , then <b>signal</b> .                 |
| r <sup>(1)</sup> | Renice a process. When prompted, enter <b>PID</b> , then <b>nice_value</b> .           |
| W                | Write (save) the current display configuration for use at the next <b>top</b> restart. |
| q                | Quit.                                                                                  |
| Note:            | <sup>(1)</sup> Not available if top started in secure mode. See <b>top(1)</b> .        |

### References

GNOME System Monitor

- o **yelp help:gnome-system-monitor**

**ps(1)**, **top(1)**, **uptime(1)**, and **w(1)** man pages

## 7.8. Practice: Monitoring Process Activity

### Guided exercise

In this lab, students will use the **top** command to dynamically view, sort, and stop processes.

### Outcomes

Practice with managing processes in real time.



### Before you begin...

Perform the following tasks as **student** on the serverX machine. Run **lab process101 setup** on serverX to prepare for this exercise.

```
[student@serverX ~]$ lab process101 setup
```

1. Open two terminal windows, side by side, to be referred to as *left* and *right*. In the right terminal, run the **top** utility. Size the window to be as tall as possible.

```
[student@serverX ~]$ top
```

2. In the left terminal, determine the number of logical CPUs on this virtual machine.

```
[student@serverX ~]$ grep "model name" /proc/cpuinfo | wc
-l
```

```
1
```

3. In the left terminal, run a single instance of the **process101** executable.

```
[student@serverX ~]$ process101
```

4. In the right terminal, observe the **top** display. Use the single keystrokes **l**, **t**, and **m** to toggle the load, threads, and memory header lines. After observing this behavior, ensure that all headers are displaying.
5. Note the process ID (PID) for **process101**. View the CPU percentage for the process, which is expected to hover around 25% or 30%.

View the load averages. On a single-CPU virtual machine, for example, the one-minute load average is currently less than a value of 1. The value observed may be affected by resource contention from another virtual machine or the virtual host.

6. In the left terminal, run a second instance of **process101**.

```
[student@serverX ~]$ process101
```

7. In **top**, note the process ID (PID) for the second **process101**. View the CPU percentage for the process, also expected to hover around 25% or 30%.

View the one-minute load average again, which may still be less than 1. Wait up to one minute to allow the calculation to adjust to the new workload.

8. In the left terminal, run a third instance of **process101**.

```
[student@serverX ~]$ process101
```

9. In **top**, note the process ID (PID) for the third **process101**. View the CPU percentage for the process, again expected to hover around 25% or 30%.

View the one-minute load average again, which now is expected to be above 1. Wait up to one minute to allow the calculation to again adjust to the new workload.

10. *Optional:* If this virtual machine has more than one logical CPU, slowly start additional **process101** instances until the one-minute load average equals or exceeds the number of logical CPUs. Divide the load average value by the number of CPUs to determine the estimated load average per CPU.
11. When finished observing the load average values, terminate each of the **process101** processes from within **top**.
  - a. Press **k**. Observe the prompt below the headers and above the columns.
  - b. Type the PID for one of the **process101** instances. Press **Enter**.
  - c. Press **Enter** again to use the default **SIGTERM** signal **15**.  
 Confirm that the selected process is no longer observed in **top**. If the PID still remains, repeat these terminating steps, substituting **SIGKILL** signal **9** when prompted.
12. Repeat the previous step for each remaining **process101** instance. Confirm that no **process101** instances remain in **top**.
13. In the right window, press **q** to exit **top**. Close extra terminal windows.

## 7.9. Lab: Monitoring and Managing Linux Processes

### Performance checklist

In this lab, students will locate and manage processes that are using the most resources on a system.

#### Outcomes

Experience using **top** as a process management tool.

#### Before you begin...

Perform the following tasks as **student** on the serverX machine. Run **lab processes setup** as **root** on serverX to prepare for this exercise.

```
[student@serverX ~]$ lab processes setup
```

1. In a terminal window, run the **top** utility. Size the window to be as tall as possible.

```
2. [student@serverX ~]$ top
3. top - 12:47:46 up 2:02, 3 users, load average: 1.67, 1.25, 0.73
4. Tasks: 361 total, 6 running, 355 sleeping, 0 stopped, 0 zombie
5. %Cpu(s): 98.5 us, 1.4 sy, 0.0 ni, 0.0 id, 0.0 wa, 0.0 hi, 0.1 si, 0.0 st
6. KiB Mem: 2043424 total, 897112 used, 1146312 free, 1740 buffers
7. KiB Swap: 4079612 total, 0 used, 4079612 free. 296276 cached Me
8.
9. PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND
10. 4019 root 20 0 4156 76 0 R 57.5 0.0 2:54.15 hippo
11. 2492 student 20 0 1359500 168420 37492 S 16.8 8.2 3:55.58 gnome-shell
12. 1938 root 20 0 189648 35972 7568 R 1.9 1.8 0:29.66 Xorg
```

```
13. 2761 student 20 0 620192 19688 12296 S 0.4 1.0 0:04.48 gnome-termi+
```

*output truncated*

14. Observe the **top** display. The default display sorts by CPU utilization, highest first. What are the processes using the most CPU time?

In addition to the default GNOME shell, find the process named **hippo**.

15. Change the display to sort by the amount of memory in use by each process.

Press **M**.

```
top - 12:57:38 up 2:11, 3 users, load average: 2.09, 1.70, 1.19
Tasks: 360 total, 5 running, 355 sleeping, 0 stopped, 0 zombie
%Cpu(s): 99.8 us, 0.2 sy, 0.0 ni, 0.0 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem: 2043424 total, 896952 used, 1146472 free, 1740 buffers
KiB Swap: 4079612 total, 0 used, 4079612 free. 296280 cached Mem
```

| PID  | USER    | PR | NI | VIRT    | RES    | SHR   | S | %CPU | %MEM | TIME+   | COMMAND      |
|------|---------|----|----|---------|--------|-------|---|------|------|---------|--------------|
| 2492 | student | 20 | 0  | 1359500 | 168420 | 37492 | S | 0.5  | 8.2  | 4:01.04 | gnome-shell  |
| 4013 | root    | 20 | 0  | 55360   | 51208  | 152   | S | 0.0  | 2.5  | 0:00.43 | elephant     |
| 1938 | root    | 20 | 0  | 189648  | 35972  | 7568  | R | 0.2  | 1.8  | 0:30.49 | Xorg         |
| 2576 | student | 20 | 0  | 533752  | 33684  | 27784 | S | 0.1  | 1.6  | 0:09.29 | vmtoolsd     |
| 2420 | student | 20 | 0  | 916268  | 25616  | 14404 | S | 0.0  | 1.3  | 0:00.61 | gnome-setti+ |
| 2550 | student | 20 | 0  | 1048204 | 23136  | 16060 | S | 0.0  | 1.1  | 0:00.46 | nautilus     |

*output truncated*

16. What are the processes with the largest memory allocations?

In addition to the default GNOME shell and **Xorg**, find a process named **elephant**.

17. Turn off the use of bold in the display. Save this configuration for reuse when **top** is restarted.

Press the single uppercase keystroke **B** to toggle bold use off.

Press the single uppercase keystroke **W** to save this configuration. The default configuration file is **.toprc** in the current user's home directory.

18. Exit **top**, then restart it again. Confirm that the new display uses the saved configuration; i.e., the display starts sorted by memory utilization and bold is turned off.

Press **q** to quit the current display, then run **top** again.

```
[student@serverX ~]$ top
```

19. Modify the display to again sort by CPU utilization. Turn on the use of bold. Observe that only *Running* or *Runnable* (state **R**) process entries are bold. Save this configuration.

Press the single uppercase keystroke **P** to sort by CPU utilization.

Press the single uppercase keystroke **B** to toggle bold use on.

Press the single uppercase keystroke **W** to save this configuration.

20. Open another terminal window if necessary. As **root**, suspend the **hippo** process. In **top**, observe that the process state is now **T**.

```
21. [student@serverX ~]$ su -
```

```
22. Password: redhat
```

```
[root@serverX ~]# pkill -SIGSTOP hippo
```

23. The **hippo** process quickly disappears from the display, since it is no longer actively using CPU resources. List the process information from the command line to confirm the process state.

```
[root@serverX ~]# ps -f $(pgrep hippo)
```

24. Resume execution of the **hippo** processes.

```
[root@serverX ~]# pkill -SIGCONT hippo
```

25. When finished observing the display, terminate the extra processes using the command line. Confirm that the processes no longer display in **top**.

```
26. [root@serverX ~]# pkill elephant
```

```
[root@serverX ~]# pkill hippo
```

27. Check that the cleanup is successful by running the grading script. If necessary, find and terminate processes listed by the grading script, and repeat grading.

```
[root@serverX ~]# lab processes grade
```

28. Exit the **top** display. Close extra terminal windows.

Press **q** to quit.

# Chapter 8. Controlling Services and Daemons

8.1. Identifying Automatically Started System Processes

8.2. Practice: Identify the Status of `systemd` Units

8.3. Controlling System Services

8.4. Practice: Using `systemctl` to Manage Services

`systemctl` 集成了 `service` 和 `chkconfig` 的功能

8.5. Lab: Controlling Services and Daemons

Introduction RHEL7使用`systemd`控制和监视网络服务和系统服务进程，代替了之前的`sysvinit`管理

| Overview   |                                                                                                                                                                                                                                          |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Goal       | To control and monitor network services and system daemons using <code>systemd</code> .                                                                                                                                                  |
| Objectives | <ul style="list-style-type: none"><li>List system daemons and network services started by the <code>systemd</code> service and socket units.</li><li>Control system daemons and network services using <code>systemctl</code>.</li></ul> |
| Sections   | <ul style="list-style-type: none"><li>Identifying Automatically Started System Processes (and Practice)</li><li>Controlling System Services (and Practice)</li></ul>                                                                     |
| Lab        | <ul style="list-style-type: none"><li>Controlling Services and Daemons</li></ul>                                                                                                                                                         |

## 8.1. Identifying Automatically Started System Processes

### Objectives

After completing this section, students should be able to list system daemons and network services started by the `systemd` service and socket units.

### Introduction to `systemd` `systemd`介绍

Viewing service status 查看服务的状态

System startup and server processes are managed by the `systemd System and Service Manager`. This program provides a method for activating system resources, server daemons, and other processes, both at boot time and on a running system.

执行不同的任务

Daemons are processes that wait or run in the background performing various tasks. Generally, daemons start automatically at boot time and continue to run until shutdown or until they are manually stopped. By convention, the names of many daemon programs end in the letter "d".

To listen for connections, a daemon uses a socket. This is the primary communication channel with local or remote clients. Sockets may be created by daemons or may be separated from the daemon and

be created by another process, such as `systemd`. The socket is passed to the daemon when a connection is established by the client.

A *service* often refers to one or more daemons, but starting or stopping a service may instead make a one-time change to the state of the system, which does not involve leaving a daemon process running afterward (called **oneshot**).

A bit of history

一次通过

服务通常是指一个或多个守护进程，但是启动或停止服务也许不是改变一次，这不包含结束进程

For many years, **process ID 1** of Linux and UNIX systems has been the **init process**. This process was **responsible for activating other services** on the system and is the origin of the term "**init system**."

Frequently used daemons were started on systems at boot time with *System V* and *LSB* init scripts.

These are shell scripts, and may vary from one distribution to another. Less frequently used daemons were started on demand by another service, such as **initd** or **xinetd**, which listens for client connections. These systems have several limitations, **which are addressed with systemd**.

In Red Hat Enterprise Linux 7, **process ID 1 is systemd**, the new init system. A few of the new features provided by `systemd` include:

- Parallelization capabilities, which increase the boot speed of a system. 并行功能，增加系统的启动速度
- On-demand starting of daemons without requiring a separate service 按需启动守护进程，不需要一个服务单独的
- Automatic service dependency management, which can prevent long timeouts, such as by not starting a network service when the network is not available. 自动服务依赖关系管理，它可以防止超时，如当网络不可用，不再启动一个网络服务
- A method of tracking related processes together by using Linux control groups. 使用Linux控制组追踪相关进程的模式

#### Note

With `systemd`, shell-based service scripts are used only for a few legacy services. Therefore, **configuration files with shell variables**, such as those found in **/etc/sysconfig**, are being replaced. Those still in use are included as `systemd` environment files and read as `NAME=VALUE` pairs. They are no longer sourced as a shell script.

#### `systemctl` and `systemd` units

The `systemctl` command is used to manage different types of `systemd` objects, called *units*. A list of available unit types can be displayed with **`systemctl -t help`**.

#### Important

省略

The `systemctl` may **abbreviate** or "ellipsize" unit names, process tree entries, and unit descriptions unless run with the `-l` option. 缩写

Some common unit types are listed below: **service 扩展名**

- **Service units** have a **.service extension** and represent system services. **This type of unit is used to start frequently accessed daemons**, such as a web server. 频繁被访问的进程
- **Socket units** have a **.socket extension** and represent **inter-process communication (IPC)** **clients**. Control of the socket will be passed to a daemon or newly started service **when a client connection is made**. Socket units are used to delay the start of a service at boot time and to start **less frequently** used services on demand. These are similar in principle to services which use the **xinetd** superserver to start on demand. 在开机启动时服务启动有延迟和很少被使用的服务
- **Path units** have a **.path extension** and are used to delay the activation of a service until a specific file system change occurs. This is commonly used for services which use `spool` directories, such as a printing system. 延迟激活一个服务，直到指定的文件系统改变发生，通常用于服务使用了 `spool` 目录，如打印服务

客户端成功建立连接

#### Service states

```
[root@rhce ~]# systemctl -t help
```

```
Available unit types:
```

```
service
```

```
socket
```

```
target
```

```
device
```

```
mount
```

```
automount
```

```
snapshot
```

```
timer
```

```
swap
```

```
path
```

```
slice
```

```
scope
```

The status of a service can be viewed with **systemctl status name.type**. If the unit type is not provided, **systemctl** will show the status of a service unit, if one exists.

```
[root@serverX ~]# systemctl status sshd.service
sshd.service - OpenSSH server daemon

 Loaded: loaded (/usr/lib/systemd/system/sshd.service; enabled)
 Active: active (running) since Thu 2014-02-27 11:51:39 EST; 7h ago
 Main PID: 1073 (sshd)
 CGroup: /system.slice/sshd.service
 └─1073 /usr/sbin/sshd -D

Feb 27 11:51:39 server0.example.com systemd[1]: Started OpenSSH server daemon.
Feb 27 11:51:39 server0.example.com sshd[1073]: Could not load host key: /et...y
Feb 27 11:51:39 server0.example.com sshd[1073]: Server listening on 0.0.0.0
Feb 27 11:51:39 server0.example.com sshd[1073]: Server listening on :: port 22.
Feb 27 11:53:21 server0.example.com sshd[1270]: error: Could not load host k...y
Feb 27 11:53:22 server0.example.com sshd[1270]: Accepted password for root f...2

Hint: Some lines were ellipsized, use -l to show in full.
```

Several keywords indicating the state of the service can be found in the status output:

| Keyword:            | Description:                                                             |
|---------------------|--------------------------------------------------------------------------|
| loaded              | Unit configuration file has been processed.                              |
| active<br>(running) | Running with one or more continuing processes.                           |
| active (exited)     | Successfully completed a one-time configuration.                         |
| active<br>(waiting) | Running but waiting for an event.                                        |
| inactive            | Not running.                                                             |
| enabled             | Will be started at boot time.                                            |
| disabled            | Will not be started at boot time.                                        |
| static              | Can not be enabled, but may be started by an enabled unit automatically. |

**Note**

The **systemctl status NAME** command replaces the **service NAME status** command used in previous versions of Red Hat Enterprise Linux.

**Listing unit files with systemctl** 取代了之前service Name status命令

In this example, please follow along with the next steps while your instructor demonstrates obtaining status information of services.

**Note**



Notice that the **systemctl** command will automatically paginate the output with **less**.

1. Query the state of all units to verify a system startup.

```
2. [root@serverX ~]# systemctl
```

查询系统启动的所有单元的状态

3. Query the state of only the service units.

```
4. [root@serverX ~]# systemctl --type=service
```

查询服务的状态

5. Investigate any units which are in a failed or maintenance state. Optionally, add the **-l** option to show the full output.

```
6. [root@serverX ~]# systemctl status rngd.service -l
```

显示服务的状态

7. The **status** argument may also be used to determine if a particular unit is active and show if the unit is enabled to start at boot time. Alternate commands can also easily show the active and enabled states:

```
8. [root@serverX ~]# systemctl is-active sshd
```

```
9. [root@serverX ~]# systemctl is-enabled sshd
```

查看服务是否开机启动

10. List the active state of all loaded units. Optionally, limit the type of unit. The **--all** option will add inactive units.

```
11. [root@serverX ~]# systemctl list-units --type=service
```

```
12. [root@serverX ~]# systemctl list-units --type=service --all
```

显示所有载入单元的active状态，all选项还会显示inactive状态

13. View the enabled and disabled settings for all units. Optionally, limit the type of unit.

```
14. [root@serverX ~]# systemctl list-unit-files --type=service
```

15. View only failed services.

仅查看错误的服务状态

```
16. [root@serverX ~]# systemctl --failed --type=service
```

## References

**systemd(1)**, **systemd.unit(5)**, **systemd.service(5)**, **systemd.socket(5)**, and **systemctl(1)** man pages

Additional information may be available in the chapter on managing services with **systemd** in the *Red Hat Enterprise Linux System Administrator's Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

## 8.2. Practice: Identify the Status of **systemd** Units

### Guided exercise

In this lab, you will identify installed and running services on the system.

Outcomes:

A list of active and enabled services on the system.

Before you begin...

Reset your serverX system.

1. List all service units on the system.

```
[student@serverX ~]$ sudo systemctl list-units
--type=service
```

2. List all socket units, active and inactive, on the system.

```
[student@serverX ~]$ sudo systemctl list-units
--type=socket --all
```

3. Explore the status of the **chronyd** service. This service is used for network time synchronization (NTP).

考察

- a. Display the status of the **chronyd** service. Note the process ID of any active daemons.

```
[student@serverX ~]$ sudo systemctl status chronyd
```

- b. Confirm that the listed daemons are running.

```
[student@serverX ~]$ ps -p PID
```

4. Explore the status of the **sshd** service. This service is used for secure encrypted communication between systems.

- a. Determine if the **sshd** service is enabled to start at system boot.

```
[student@serverX ~]$ sudo systemctl is-enabled
sshd
```

- b. Determine if the **sshd** service is active without displaying all of the status information.

```
[student@serverX ~]$ sudo systemctl is-active sshd
```

- c. Display the status of the **sshd** service.

```
[student@serverX ~]$ sudo systemctl status sshd
```

5. List the enabled or disabled states of all service units.

```
[student@serverX ~]$ sudo systemctl list-unit-files
--type=service
```

## 8.3. Controlling System Services

### Objectives

After completing this section, students should be able to **control system daemons and network services using `systemctl`**.

使用systemctl控制系统进程和网络服务

### Starting and stopping system daemons on a running system

Controlling system daemons

对于一个服务，当更改了配置文件或做了其他更新时，需要重启服务

Changes to a configuration file or other updates to a service may require that the service be restarted.

A service that is no longer used may be stopped before removing the software. A service that is not frequently used may be manually started by an administrator only when it is needed.

在软件被移除之前如果长时间不使用可以把服务停止

一个服务很少使用的话可以配置成手动启动

In this example, please follow along with the next steps while your instructor demonstrates managing services on a running system.

1. View the status of a service.

```
2. [root@serverX ~]# systemctl status sshd.service
```

3. Verify that the process is running.

```
4. [root@serverX ~]# ps -up PID
```

5. Stop the service and verify the status.

```
6. [root@serverX ~]# systemctl stop sshd.service
```

```
7. [root@serverX ~]# systemctl status sshd.service
```

8. Start the service and view the status. The process ID has changed.

```
9. [root@serverX ~]# systemctl start sshd.service
```

```
10. [root@serverX ~]# systemctl status sshd.service
```

11. Stop, then start, the service in a single command.

```
12. [root@serverX ~]# systemctl restart sshd.service
```

```
13. [root@serverX ~]# systemctl status sshd.service
```

14. Issue instructions for a service to read and reload its configuration file without a complete stop and start. The process ID will not change.

重新读取和载入配置文件不需要关闭和启动服务，可以使用reload，进程号不会发生改变

```
15. [root@serverX ~]# systemctl reload sshd.service
```

```
16. [root@serverX ~]# systemctl status sshd.service
```

Unit dependencies

服务启动时被依赖于其他服务，如一个网络连接或文件系统触发

Services may be started as dependencies of other services. If a socket unit is enabled and the service unit with the same name is not, the service will automatically be started when a request is made on the network socket. Services may also be triggered by path units when a file system condition is met. For example, a file placed into the print spool directory will cause the **cups** print service to be started if it is not running.

```
[root@serverX ~]# systemctl stop cups.service
Warning: Stopping cups, but it can still be activated by:
 cups.path
 cups.socket
```

To completely stop printing services on a system, stop all three units. Disabling the service will disable the dependencies.

The **systemctl list-dependencies UNIT** command can be used to print out a tree of what other units must be started if the specified unit is started. Depending on the exact dependency, the other unit may need to be running before or after the specified unit starts.

The **--reverse** option to this command will show what units need to have the specified unit started in order to run.

#### Masking services

有时候，系统安装了多个服务，可能会产生冲突

At times, a system may have conflicting services installed. For example, there are multiple methods to manage networks (network and NetworkManager) and firewalls (iptables and firewalld). To prevent an administrator from accidentally starting a service, that service may be *masked*. Masking will create a link in the configuration directories so that if the service is started, nothing will happen.

意外地

```
[root@serverX ~]# systemctl mask network
ln -s '/dev/null' '/etc/systemd/system/network.service'
[root@serverX ~]# systemctl unmask network
rm '/etc/systemd/system/network.service'
```

如果一个服务被伪装了，会创建一个链接在配置目录，如果服务被开启了，它什么也不会做

#### Important

A disabled service will not be started automatically at boot or by other unit files, but can be started manually. A masked service can not be started manually or automatically.

保证

## Enabling system daemons to start or stop at boot

Starting a service on a running system does not guarantee that the service will be started when the system reboots. Similarly, stopping a service on a running system will not keep it from starting again when the system reboots. Services are started at boot time when links are created in the appropriate **systemd** configuration directories. These links are created and removed with **systemctl** commands.

设置服务开机启动会在systemd配置目录创建合适的链接，可以使用systemctl工具配置

In this example, please follow along with the next steps while your instructor demonstrates enabling and disabling services.

1. View the status of a service.

```
2. [root@serverX ~]# systemctl status sshd.service
```

3. Disable the service and verify the status. Note that disabling a service does not stop the service.

```
4. [root@serverX ~]# systemctl disable sshd.service
5. [root@serverX ~]# systemctl status sshd.service
```

6. Enable the service and verify the status.

```
7. [root@serverX ~]# systemctl enable sshd.service
8. [root@serverX ~]# systemctl is-enabled sshd.service
```

## Summary of `systemctl` commands

Services can be started and stopped on a running system and enabled or disabled for automatic start at boot time.

| Task:                                                                       | Command:                                      |
|-----------------------------------------------------------------------------|-----------------------------------------------|
| View detailed information about a unit state.                               | <code>systemctl status UNIT</code>            |
| Stop a service on a running system.                                         | <code>systemctl stop UNIT</code>              |
| Start a service on a running system.                                        | <code>systemctl start UNIT</code>             |
| Restart a service on a running system.                                      | <code>systemctl restart UNIT</code>           |
| Reload configuration file of a running service.                             | <code>systemctl reload UNIT</code>            |
| Completely disable a service from being started, both manually and at boot. | <code>systemctl mask UNIT</code>              |
| Make a masked service available.                                            | <code>systemctl unmask UNIT</code>            |
| Configure a service to start at boot time.                                  | <code>systemctl enable UNIT</code>            |
| Disable a service from starting at boot time.                               | <code>systemctl disable UNIT</code>           |
| List units which are required and wanted by the specified unit.             | <code>systemctl list-dependencies UNIT</code> |

### References

`systemd(1)`, `systemd.unit(5)`, `systemd.service(5)`, `systemd.socket(5)`, and `systemctl(1)` man pages

Additional information may be available in the chapter on managing services with `systemd` in the *Red Hat Enterprise Linux System Administrator's Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

## 8.4. Practice: Using `systemctl` to Manage Services

### Guided exercise

In this lab, you will manage a service unit that is already installed on the system.

#### Outcomes:

The `chronyd` service is disabled and no longer running on the system.

#### Before you begin...

Reset your serverX system.

1. Observe the results of **systemctl restart** and **systemctl reload** commands.

- a. Display the status of the **sshd** service. Note the process ID of the daemon.

```
[student@serverX ~]$ sudo systemctl status sshd
```

- b. Restart the **sshd** service and view the status. The process ID of the daemon has changed.

- c. 

```
[student@serverX ~]$ sudo systemctl restart sshd
```

```
[student@serverX ~]$ sudo systemctl status sshd
```

- d. Reload the **sshd** service and view the status. The process ID of the daemon has not changed and connections have not been interrupted.

- e. 

```
[student@serverX ~]$ sudo systemctl reload sshd
```

```
[student@serverX ~]$ sudo systemctl status sshd
```

2. Verify that the **chronyd** service is running.

```
[student@serverX ~]$ sudo systemctl status chronyd
```

3. Stop the **chronyd** service and view the status.

4. 

```
[student@serverX ~]$ sudo systemctl stop chronyd
```

```
[student@serverX ~]$ sudo systemctl status chronyd
```

5. Determine if the **chronyd** service is enabled to start at system boot.

```
[student@serverX ~]$ sudo systemctl is-enabled chronyd
```

6. Reboot the system, then view the status of the **chronyd** service.

```
[student@serverX ~]$ sudo systemctl status chronyd
```

7. Disable the **chronyd** service so that it does not start at system boot, then view the status of the service.

8. 

```
[student@serverX ~]$ sudo systemctl disable chronyd
```

```
[student@serverX ~]$ sudo systemctl status chronyd
```

9. Reboot the system, then view the status of the **chronyd** service.

```
[student@serverX ~]$ sudo systemctl status chronyd
```

## ▮ 8.5. Lab: Controlling Services and Daemons

### Performance checklist

In this lab, you will manage a service unit that is already installed on the system.

#### Outcomes:

The **psacct** service is enabled and running on the system, and the **rsyslog** service is disabled and no longer running on the system.

#### Before you begin...

Reset your serverX system.

1. Start the **psacct** service.

```
2. [student@serverX ~]$ sudo systemctl start psacct
```

```
[student@serverX ~]$ sudo systemctl status psacct
```

3. Configure the **psacct** service so that it starts at system boot.

```
4. [student@serverX ~]$ sudo systemctl enable psacct
```

```
[student@serverX ~]$ sudo systemctl status psacct
```

5. Stop the **rsyslog** service.

```
6. [student@serverX ~]$ sudo systemctl stop rsyslog
```

```
[student@serverX ~]$ sudo systemctl status rsyslog
```

7. Configure the **rsyslog** service so that it does not start at system boot.

```
8. [student@serverX ~]$ sudo systemctl disable rsyslog
```

```
[student@serverX ~]$ sudo systemctl status rsyslog
```

9. Reboot the system, then run **lab services grade** to verify the configuration.

```
[student@serverX ~]$ lab services grade
```

# ▮ Chapter 9. Configuring and Securing OpenSSH Service

## 9.1. Accessing the Remote Command Line with SSH

9.2. Practice: Accessing the Remote Command Line

9.3. Configuring **SSH Key-based** Authentication

9.4. Practice: Using SSH Key-based Authentication

9.5. Customizing SSH Service Configuration

9.6. Practice: Restricting SSH Logins

9.7. Lab: Configuring and Securing OpenSSH Service

## Introduction

| Overview   |                                                                                                                                                                                                                                                                                                                                                                                        |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Goal       | To configure secure command-line access on remote systems using OpenSSH.                                                                                                                                                                                                                                                                                                               |
| Objectives | <ul style="list-style-type: none"><li>Log into a remote system using <code>ssh</code> to run commands from a shell prompt.</li><li>Set up <code>ssh</code> to allow secure password-free logins by using a private authentication key file.</li><li>Customize <code>sshd</code> configuration to restrict direct logins as root or to disable password-based authentication.</li></ul> |
| Sections   | <ul style="list-style-type: none"><li>Accessing the Remote Command Line with SSH (and Practice)</li><li>Configuring SSH Key-based Authentication (and Practice)</li><li>Customizing SSH Service Configuration (and Practice)</li></ul>                                                                                                                                                 |
| Lab        | <ul style="list-style-type: none"><li>Configuring and Securing OpenSSH Service</li></ul>                                                                                                                                                                                                                                                                                               |

## 9.1. Accessing the Remote Command Line with SSH

### Objective

After completing this section, students should be able to **log into a remote system using `ssh`** to run commands from a shell prompt.

通过SSH远程登录主机

*Accessing the command line using `ssh`*

### What is the OpenSSH **secure shell (SSH)**?

在远程系统上安全的使用shell

The term OpenSSH refers to the software implementation of the **Secure Shell** software used in the system. The OpenSSH **Secure Shell**, `ssh`, is used to **securely run a shell on a remote system**. If you have a user account on a remote Linux system providing SSH services, `ssh` is the command normally used to remotely log into that system. **The `ssh` command can also be used to run an individual command on a remote system.**

在远程系统上也可以通过ssh独立地调用命令

### Secure Shell examples

Here are some examples of `ssh` command syntax for remote login and remote execution:

- Create a remote interactive shell as the current user, then return to your previous shell when done with the `exit` command.



- `[student@host ~]$ ssh remotehost`
- `student@remotehost's password:`
- `[student@remotehost ~]$ exit`
- `Connection to remotehost closed.`

```
[student@host ~]$
```

- Connect to a remote shell as a different user (**remoteuser**) on a selected host (**remotehost**):

- `[student@host ~]$ ssh remoteuser@remotehost`
- `remoteuser@remotehost's password:`

```
[remoteuser@remotehost ~]$
```

- Execute a single command (**hostname**) on a remote host (**remotehost**) and as a remote user (**remoteuser**) in a way that returns the output to the local display:

- `[student@host ~]$ ssh remoteuser@remotehost hostname`
- `remoteuser@remotehost's password:`
- `remotehost.example.com`

```
[student@host ~]$
```

特别有用，对于查看ssh登录来说

The **w** command displays a list of users currently logged into the computer. This is especially useful to show which users are logged in using **ssh** from which remote locations, and what they are doing.

```
[student@host ~]$ w -f
USER TTY FROM LOGIN@ IDLE JCPU PCPU WHAT
student tty1 :0 Wed08 2days 1:52m 0.07s pam: gdm-passwo
root tty6 - 12:33 4:14m 16.27s 15.74s -bash
student pts/0 :0.0 Wed08 5:11 1.63s 1.60s /usr/bin/gnome-
student pts/1 :0.0 Wed08 43:44 14.48s 13.81s vim hello.c
student pts/3 :0.0 Wed14 0.00s 0.06s 0.06s w
visitor pts/6 server2.example. 09:22 3:14 0.02s 0.02s -bash
```

In this example, user *student* logged in on virtual console 1 (**tty1**) through the graphical login (**:0**) at about 08:00 on Wednesday. User *student* currently has three pseudo-terminals open (**pts/0**, **pts/1**, and **pts/3**) started by the graphical environment; these are almost certainly terminal windows. In one window, *student* is editing **hello.c**. User *root* is logged in on virtual console 6, starting at 12:33 today. User *visitor* logged in on pseudo-terminal 6 at 09:22 today from the host `server2.example.com` (note that the name has been truncated), probably using **ssh**, and has been sitting idle at a shell prompt for three minutes and 14 seconds.

## SSH host keys

公钥加密

SSH secures communication through **public-key encryption**. When an **ssh** client connects to an SSH server, before the client logs in, **the server sends it a copy of its public key**. This is used to set up the secure encryption for the communication channel and to authenticate the server to the client.

The first time a user uses **ssh** to connect to a particular server, the **ssh** command stores the server's public key in the user's **~/.ssh/known\_hosts** file. Every time the user connects after that, the client makes sure it gets the same public key from the server by comparing the server's entry in the **~/.ssh/known\_hosts** file to the public key the server sent. If the keys do *not* match, the client assumes that the **network traffic is being hijacked** or that the server has been compromised, and **breaks the connection**.

阻止连接

网络流量被劫持

服务器遭到泄露

This means that if a server's public key is changed (because the key was lost due to hard drive failure, or replaced for some legitimate reason), users will need to update

their **~/.ssh/known\_hosts** files and remove the old entry in order to log in.

- Host IDs are stored in **~/.ssh/known\_hosts** on your local client system:

```
○ $ cat ~/.ssh/known_hosts
```

```
remotehost,192.168.0.101 ssh-rsa AAAAB3Nzac...
```

- Host keys are stored in **/etc/ssh/ssh\_host\_key\*** on the SSH server.

```
○ $ ls /etc/ssh/*key*
```

```
○ ssh_host_dsa_key ssh_host_key ssh_host_rsa_key
```

```
ssh_host_dsa_key.pub ssh_host_key.pub ssh_host_rsa_key.pub
```

### Note

An even better approach is to add entries matching a server's **ssh\_host\_\*key.pub** files to user **~/.ssh/known\_hosts** or the systemwide **/etc/ssh/ssh\_known\_hosts** in advance when the public keys change. See **ssh-copy-id(1)** for an advanced way to manage SSH keys.

### References

Additional information may be available in the chapter on using the **ssh** utility in the *Red Hat Enterprise Linux System Administrator's Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

**ssh(1)**, **w(1)**, and **hostname(1)** man pages

## 9.2. Practice: Accessing the Remote Command Line

### Guided exercise

In this lab, students will log into a remote system as different users and execute commands.

#### Outcomes:

Students will log into a remote system and execute commands with the OpenSSH secure shell.

1. Log in as student on your desktopX machine.
2. **ssh** to your serverX machine. Accept the host key if asked. The host key is recorded on our local machine to identify the remote machine. The **ssh** command will fail to execute

properly if the remote ssh host appears to have a different key than the recorded host key. The host key records are stored in the **known\_hosts** file in the **.ssh** directory in the user's home directory on the local system.

```
3. [student@desktopX ~]$ ssh student@serverX
4. The authenticity of host 'serverX (172.25.X.11)' can't be established.
5. ECDSA key fingerprint is 47:bf:82:cd:fa:68:06:ee:d8:83:03:1a:bb:29:14:a3.
6. Are you sure you want to continue connecting (yes/no)? yes
```

```
student@serverX's password: student
```

7. Run the **w** command. The output of the **w** clearly indicates we have logged in as user student from desktopX.

```
8. [student@serverX ~]$ w -f
9. 11:01:23 up 1 day, 19:10, 1 user, load average: 0,0,0
10. USER TTY FROM LOGIN@ IDLE JCPU PCPU WHAT
```

```
student pts/1 desktopX 11:01 0.00s 0.12s 0.09s w
```

11. Execute the **exit** command to terminate the secure shell connection.

```
12. [student@serverX ~]$ exit
```

```
[student@desktopX ~]$
```

13. This time, **ssh** to your serverX machine as user **root**.

```
14. [student@desktopX ~]$ ssh root@serverX
15. root@serverX's password: redhat
```

```
[root@serverX ~]#
```

16. Run the **w** command again. This time, the output of the **w** shows the active connection to the root user account from desktopX.

```
17. [root@serverX ~]# w -f
18. 11:01:23 up 1 day, 19:10, 1 user, load average: 0,0,0
19. USER TTY FROM LOGIN@ IDLE JCPU PCPU
```

```
root pts/2 desktopX 11:09 0.00s 0.13s 0.08s w
```

20. Run the **exit** to terminate the secure shell connection.

```
21. [root@serverX ~]# exit
```

```
[student@desktopX ~]$
```

22. There are different reasons why a remote host might have legitimately changed its host key. One common reason is when the remote machine is replaced because of hardware failure, or reinstalled. Usually, it is advisable to only remove the key entry for the particular host in the **known\_hosts**. In this case, there is only one host entry in the **known\_hosts**, so it can be removed completely. Remove the **known\_hosts** file for the user student.

```
[student@desktopX ~]$ rm ~/.ssh/known_hosts
```

23. **ssh** to serverX as **root** again. Accept the key, log in, and then exit the session.

```
24. [student@desktopX ~]$ ssh root@serverX
25. The authenticity of host 'serverX (:::1)' can't be established.
26. ECDSA key fingerprint is 47:bf:82:cd:fa:68:06:ee:d8:83:03:1a:bb:29:14:a3.
27. Are you sure you want to continue connecting (yes/no)? yes
28. root@serverX's password: redhat
29. [root@serverX ~]# exit
```

```
[student@desktopX ~]$
```

30. Use **ssh** non-interactively to run the **hostname** command on serverX as **root**.

```
31. [student@desktopX ~]$ ssh root@serverX hostname
32. root@serverX's password: redhat
```

```
serverX.example.com
```

## 9.3. Configuring SSH Key-based Authentication

基于密钥认证

### Objective

After completing this section, students should be able to set up SSH to allow secure logins without passwords by using a private authentication key file.

Using SSH keys

### SSH key-based authentication

Users can authenticate **ssh** logins without a password by using *public key authentication*. **ssh** allows users to authenticate using a private-public key scheme. This means that two keys are generated, a private key and a public key. The private key file is used as the authentication credential, and like a password, must be kept secret and secure. The public key is copied to systems the user wants to log into, and is used to verify the private key. The public key does not need to be secret. An SSH server that has the public key can issue a challenge that can only be answered by a system holding your private key. As a result, you can authenticate using the presence of

公钥拷贝到要登录的系统上，用于校验私钥

私钥类似于我们的口令，要安全地存放

your key. This allows you to access systems in a way that doesn't require typing a password every time, but is still secure.

Key generation is done using the `ssh-keygen` command. This generates the private key `~/.ssh/id_rsa` and the public key `~/.ssh/id_rsa.pub`.

#### Note

During key generation, there is the option to specify a passphrase which must be provided in order to access your private key. In the event the private key is stolen, it is very difficult for someone other than the issuer to use it when protected with a passphrase. This adds enough of a time buffer to make a new key pair and remove all references to the old keys before the private key can be used by an attacker who has cracked it.

It is always wise to passphrase-protect the private key since the key allows access to other machines. However, this means the passphrase must be entered whenever the key is used, making the authentication process no longer password-less. This can be avoided using `ssh-agent`, which can be given your passphrase once at the start of the session (using `ssh-add`), so it can provide the passphrase as needed while you stay logged in.

For additional information on the `ssh-agent` command, consult the Red Hat System Administration Guide, Chapter 8.2.4.2.: Configuring ssh-agent.

Once the SSH keys have been generated, they are stored by default in the `.ssh/` directory of your home directory. Permissions should be 600 on the private key and 644 on the public key.

Before key-based authentication can be used, the public key needs to be copied to the destination system. This can be done with `ssh-copy-id`.

```
[student@desktopX ~]$ ssh-copy-id root@desktopY
```

When the key is copied to another system using `ssh-copy-id`, it copies the `~/.ssh/id_rsa.pub` file by default.

#### SSH key demonstration

- Use `ssh-keygen` to create a public-private key pair.

```
[student@desktopX ~]$ 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): redhat
Enter same passphrase again: redhat
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:
a4:49:cf:fb:ac:ab:c8:ce:45:33:f2:ad:69:7b:d2:5a student@desktopX.example.com
The key's randomart image is:
+--[RSA 2048]-----+
| |
```

```

○ | |
○ | . . |
○ | . * |
○ | . * S |
○ | + + . |
○ | o.E |
○ | o oo+oo |
○ | .=.**ooo |

```

```
+-----+
```

- Use **ssh-copy-id** to copy the public key to the correct location on a remote system. For example:

```
[student@desktopX ~]$ ssh-copy-id -i ~/.ssh/id_rsa.pub
root@serverX.example.com
```

## References

Additional information may be available in the chapter on using key-based authentication in the *Red Hat Enterprise Linux System Administrator's Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

**ssh-keygen**(1), **ssh-copy-id**(1), **ssh-agent**(1), **ssh-add**(1) man pages

## 9.4. Practice: Using SSH Key-based Authentication

### Guided exercise

In this lab, you will set up SSH key-based authentication.

#### Outcomes:

Students will set up SSH user key-based authentication to initiate SSH connections.

1. Create an SSH key pair as **student** on desktopX using no passphrase.

```

2. [student@desktopX ~]$ ssh-keygen
3. Generating public/private rsa key pair.
4. Enter file in which to save the key (/home/student/.ssh/id_rsa): Enter
5. Created directory '/home/student/.ssh'.
6. Enter passphrase (empty for no passphrase): Enter
7. Enter same passphrase again: Enter
8. Your identification has been saved in /home/student/.ssh/id_rsa.
9. Your public key has been saved in /home/student/.ssh/id_rsa.pub.

```

```
...
```

10. Send the SSH public key to the **student** account on serverX.

```
11. [student@desktopX ~]$ ssh-copy-id serverX
12. The authenticity of host 'serverX (172.25.X.11)' can't be established.
13. ECDSA key fingerprint is 33:fa:a1:3c:98:30:ff:f6:d4:99:00:4e:7f:84:3e:c3.
14. Are you sure you want to continue connecting (yes/no)? yes
15. /usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter
 out any that are already installed
16. /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
17. student@serverX's password: student
18.
19. Number of key(s) added: 1
20.
21. Now try logging into the machine, with: "ssh 'student@serverX'"
22. and check to make sure that only the key(s) you wanted were added.
```

23. Run the **hostname** command by using **ssh** to display the hostname of the serverX.example.com machine without the need to enter a password.

```
24. [student@desktopX ~]$ ssh serverX 'hostname'
```

```
serverX.example.com
```

## 9.5. Customizing SSH Service Configuration

### Objective

After completing this section, students should be able to customize sshd configuration to restrict direct logins as root or to disable password-based authentication.

*Securing SSH access*

### The OpenSSH server configuration file

While OpenSSH server configuration usually does not require modification, additional security measures are available.

Various aspects of the OpenSSH server can be modified in the configuration file **/etc/ssh/sshd\_config**.

阻止root用户使用ssh登录

### Prohibit the root user from logging in using SSH

From a security standpoint, it is advisable to prohibit the root user from directly logging into the system with **ssh**.

- The username root exists on every Linux system by default, so a potential attacker only has to guess the password, instead of a valid username and password combination.
- The root user has unrestricted privileges.

The OpenSSH server has an internal configuration file setting to prohibit a system login as user root, which is commented out by default in the `/etc/ssh/sshd_config` file:

```
#PermitRootLogin yes
```

By enabling the previous option in the `/etc/ssh/sshd_config` configuration file as follows, the root user will be unable to log into the system using the `ssh` command after the `sshd` service has been restarted:

```
PermitRootLogin no
```

The `sshd` service has to be restarted to put the changes into effect:

```
[root@serverX ~]# systemctl restart sshd
```

Another option is to only allow key-based ssh login as root with:

```
PermitRootLogin without-password
```

## Prohibit password authentication using SSH

Only allowing key-based logins to the remote command line has various advantages:

- SSH keys are longer than an average password, which adds security.
- Less effort to initiate remote shell access after the initial setup.

There is an option in the `/etc/ssh/sshd_config` configuration file which turns on password authentication by default:

```
PasswordAuthentication yes
```

To prevent password authentication, the `PasswordAuthentication` option has to be set to `no` and the `sshd` service needs to be restarted:

```
PasswordAuthentication no
```

Keep in mind that whenever you change the `/etc/ssh/sshd_config` file, the `sshd` service has to be restarted:

```
[root@serverX ~]# systemctl restart sshd
```

### References

`ssh(1)`, `sshd_config(5)` man pages

## 9.6. Practice: Restricting SSH Logins

### Guided exercise

In this lab, you will enable additional security features in OpenSSH.

#### Outcomes:

Prohibit direct SSH login as root on serverX; prohibit users from using passwords to login through SSH to serverX; public key authentication should still be allowed for regular users.

Before you begin...



Reset the desktopX and serverX systems.

Run **lab ssh setup** on both desktopX and serverX. This will create a user account called **visitor** with a password of **password**.

```
[student@desktopX ~]$ lab ssh setup
```

```
[student@serverX ~]$ lab ssh setup
```

1. Generate SSH keys on desktopX, copy the public key to the **student** account on serverX, and verify that the keys are working.

- a. Generate the SSH keys on desktopX.

```
b. [student@desktopX ~]$ ssh-keygen
c. Generating public/private rsa key pair.
d. Enter file in which to save the key (/home/student/.ssh/id_rsa):
e. Enter
f.
g. Created directory '/home/student/.ssh'.
h. Enter passphrase (empty for no passphrase):
i. Enter
j.
k. Enter same passphrase again:
l. Enter
m.
n. Your identification has been saved in /home/student/.ssh/id_rsa.
o. Your public key has been saved in /home/student/.ssh/id_rsa.pub.
p. ...
```

- q. Copy the SSH public key to the **student** account on serverX.

```
r. [student@desktopX ~]$ ssh-copy-id serverX
s. The authenticity of host 'serverX (172.25.X.11)' can't be established.
t. ECDSA key fingerprint is
 33:fa:a1:3c:98:30:ff:f6:d4:99:00:4e:7f:84:3e:c3.
u. Are you sure you want to continue connecting (yes/no)? yes
v.
w. /usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s),
 to filter out any that are already installed
x. /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
```

```
y. student@serverX's password: student
z.
aa. Number of key(s) added: 1
bb.
cc. Now try logging into the machine, with: "ssh 'student@serverX'"
dd. and check to make sure that only the key(s) you wanted were added.
```

ee. Verify that key-based SSH authentication is working for user student on serverX.

```
ff. [student@desktopX ~]$ ssh student@serverX
gg. [student@serverX ~]$
```

2. Log into the serverX machine and obtain superuser privileges.

```
3. [student@desktopX ~]$ ssh student@serverX
4. [student@serverX ~]$ su -
5. Password: redhat
6. [root@serverX ~]#
```

7. Configure SSH on serverX to prevent root logins.

a. As user root, edit `/etc/ssh/sshd_config` on serverX so that "PermitRootLogin" is uncommented and set to "no."

```
PermitRootLogin no
```

b. Restart the SSH service on the serverX machine.

```
[root@serverX ~]# systemctl restart sshd
```

c. Confirm that **root** cannot log in with SSH, but **student** is permitted to log in.

```
d. [student@desktopX ~]$ ssh root@serverX
e. Password: redhat
f. Permission denied, please try again.
g. Password: redhat
h. Permission denied, please try again.
i. Password: redhat
j. Permission denied (publickey,gssapi-keyex,gssapi-with-mic,password)
k.
l. [student@desktopX ~]$ ssh student@serverX
```

```
[student@serverX ~]$
```

8. Configure SSH on serverX to prevent password authentication.
  - a. Edit the configuration file `/etc/ssh/sshd_config` as user root so that the "PasswordAuthentication" entry is set to "no":

```
PasswordAuthentication no
```

- b. Restart the SSH service.

```
[root@serverX ~]# systemctl restart sshd
```

- c. Confirm that **visitor** cannot log in using a password, but **student** is permitted to log in using the SSH keys created earlier.

```
d. [student@desktopX ~]$ ssh visitor@serverX
e. Permission denied (publickey,gssapi-keyex,gssapi-with-mic).
f. [student@desktopX ~]$ ssh student@serverX
```

```
[student@serverX ~]$
```

## 9.7. Lab: Configuring and Securing OpenSSH Service

### Performance checklist

In this lab, you will add security measures to the ssh service.

#### Outcomes:

Students will set up SSH keys, configure and exclusively allow user key-based authentication, and lock down the OpenSSH service to prevent the root user from logging into the system by using SSH.

#### Before you begin...

Reset the desktopX and serverX systems.

Run **lab ssh setup** as the **student** user on both desktopX and serverX. This will create a user account called **visitor** with a password of **password**.

```
[student@desktopX ~]$ lab ssh setup
[student@serverX ~]$ lab ssh setup
```

Unless specified, all steps are to be performed as user **visitor**.

1. Generate SSH keys on desktopX for user visitor and copy the public key to the **visitor** account on serverX.

- a. Generate a SSH public key on desktopX as user visitor.

```
b. [visitor@desktopX ~]$ ssh-keygen
```

- c. Install the SSH public key generated previously on desktopX to the **visitor** account on serverX.

```
d. [visitor@desktopX ~]$ ssh-copy-id serverX
```

```
e. The authenticity of host 'serverX (172.25.X.11)' can't be established.
f. ECDSA key fingerprint is
 xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx.
g. Are you sure you want to continue connecting (yes/no)? yes
h. /usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s),
 to filter out any that are already installed
i. /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
j. visitor@serverX's password: password
k.
l. Number of key(s) added: 1
m.
n. Now try logging into the machine, with: "ssh 'visitor@serverX'"
o. and check to make sure that only the key(s) you wanted were added.
p.
```

2. Disable ssh login for the root user and password-based SSH authentication on serverX.

a. Log into the serverX virtual machine as user root.

```
[visitor@desktopX ~]$ ssh root@serverX
```

b. Customize the ssh service on serverX by disabling SSH connections for the user root and only allow key-based login.

Set the necessary config file parameters in `/etc/ssh/sshd_config`:

```
PermitRootLogin no
PasswordAuthentication no
```

c. Restart the sshd service on serverX.

```
[root@serverX ~]# systemctl restart sshd
```

3. Verify that user root is not allowed to login to serverX by using ssh, while user visitor is with the private key.

a. On a different terminal window on desktopX, validate that user root cannot connect to serverX with the `ssh` command. It should fail because we disabled root logins with the ssh service.

```
b. [visitor@desktopX ~]$ ssh root@serverX
```

```
c. Permission denied (publickey,gssapi-keyex,gssapi-with-mic).
```

d. Try logging in as user `student` to `serverX` from `desktopX` by using `ssh`. It should fail because we did not add the public key from that user to the `student` account on the `serverX` machine.

```
e. [visitor@desktopX ~]$ ssh student@serverX
```

```
f. Permission denied (publickey,gssapi-keyex,gssapi-with-mic).
```

g. Verify the `ssh` service is still accepting key-based authentication by successfully connecting to `serverX` as user `visitor` with the `ssh` command.

```
h. [visitor@desktopX ~]$ ssh visitor@serverX
```

```
[visitor@serverX ~]$
```

# Chapter 10. Analyzing and Storing Logs

分析和存储日志

10.1. System Log Architecture

10.2. Practice: System Logging Components

10.3. Reviewing Syslog Files

10.4. Practice: Finding Log Entries

10.5. Reviewing systemd Journal Entries

10.6. Practice: Finding Events With `journalctl`

10.7. Preserving the systemd Journal

10.8. Practice: Configure a Persistent systemd Journal

10.9. Maintaining Accurate Time

精确的

10.10. Practice: Adjusting System Time

10.11. Lab: Analyzing and Storing Logs

Introduction

Overview:

|      |                                                                  |
|------|------------------------------------------------------------------|
| Goal | To locate and accurately interpret relevant system log files for |
|------|------------------------------------------------------------------|

准确定位和精确说明相关的系统日志文件对于排错非常重要

| Overview:         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                   | <b>troubleshooting</b> purposes. <b>目的</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Objectives</b> | <ul style="list-style-type: none"> <li>Describe the basic syslog architecture in Red Hat Enterprise Linux 7.</li> <li>Interpret entries in relevant syslog files to troubleshoot problems or <b>review system status.</b> <b>查看系统状态</b></li> <li>Find and interpret log entries in the systemd journal to troubleshoot problems or review system status.</li> <li>Configure <b>systemd-journald</b> to store its journal on disk rather than in memory.</li> <li>Maintain accurate time synchronization and time zone configuration to ensure correct timestamps in system logs.</li> </ul> |
| <b>Sections</b>   | <ul style="list-style-type: none"> <li>System Log Architecture (and Practice)</li> <li>Reviewing Syslog Files (and Practice)</li> <li>Reviewing systemd Journal Entries (and Practice)</li> <li>Preserving the systemd Journal (and Practice)</li> <li>Maintaining Accurate Time (and Practice)</li> </ul>                                                                                                                                                                                                                                                                                  |
| <b>Lab</b>        | <ul style="list-style-type: none"> <li>Analyzing and Storing Logs</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## 10.1. System Log Architecture

### Objectives

After completing this section, students should be able to describe the basic syslog architecture in Red Hat Enterprise Linux 7.

#### System logging

**系统日志**

**日志对于系统审计和排错非常的有用**

Processes and the operating system kernel need to be able to record a log of events that happen. These logs can be useful for **auditing the system and troubleshooting problems.** By convention, the **/var/log** directory is where these logs are persistently stored. **默认的日志存放目录**

A standard logging system based on the **syslog protocol** is built into Red Hat Enterprise Linux. **Many** **programs use this system to record events and organize them into log files.** In Red Hat Enterprise Linux

**许多应用程序也使用该系统记录他们的日志**

7, syslog messages are handled by two services, **systemd-journald** and **rsyslog.**

The **systemd-journald** daemon provides an improved log management service that collects **messages from the kernel,** the early stages of the boot process, standard output and error of daemons as they start up and run, and syslog. It writes these messages to a structured journal of events that, by default, **does not persist between reboots.** This allows syslog messages and events which are missed by syslog to be collected in one central database. **The syslog messages are also forwarded**

**by systemd-journald to rsyslog** for further processing.

**The rsyslog** service then sorts the syslog messages by type (or facility) and priority, and writes them to persistent files in the **/var/log** directory. **rsyslog服务按消息的类型和优先级排序**

The **/var/log** directory holds various system- and service-specific log files maintained by **rsyslog.**

**systemd-journald 服务提供了更优的日志管理机制**

**首先，引导进程阶段的信息写入到日志文件系统，重启时不保留**

□

## Overview of system log files

| Log file                                                                                        | Purpose                                                                                                                                                                                                                                                                    |
|-------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>/var/log/messages</code><br><span style="border: 1px solid red; padding: 2px;">周期性</span> | Most syslog messages are logged here. The exceptions are messages related to authentication and email processing, that periodically run jobs, and those which are purely debugging-related.<br><span style="border: 1px solid red; padding: 2px;">邮件和认证进程、周期性任务计划除外</span> |
| <code>/var/log/secure</code>                                                                    | The log file for security and authentication-related messages and errors.                                                                                                                                                                                                  |
| <code>/var/log/maillog</code>                                                                   | The log file with mail server-related messages.                                                                                                                                                                                                                            |
| <code>/var/log/cron</code>                                                                      | The log file related to periodically executed tasks.                                                                                                                                                                                                                       |
| <code>/var/log/boot.log</code>                                                                  | Messages related to system startup are logged here.                                                                                                                                                                                                                        |

### References

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

Additional information may be available in the *Red Hat Enterprise Linux System Administrator's Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

## □ 10.2. Practice: System Logging Components

### Quiz

Match the following items to their counterparts in the table.

| Purpose                                                                                                                                                                                     | Log file                       |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------|
| Most syslog messages are logged here. The exceptions are messages related to authentication, email processing, and that periodically run jobs, or those which are purely debugging-related. | <code>/var/log/messages</code> |
| The log file for security and authentication-related messages and errors.                                                                                                                   | <code>/var/log/secure</code>   |
| The directory to which <code>rsyslog</code> is writing all the log files.                                                                                                                   | <code>/var/log</code>          |
| The log file with mail server-related messages.                                                                                                                                             | <code>/var/log/maillog</code>  |
| The log file related to periodically executed tasks.                                                                                                                                        | <code>/var/log/cron</code>     |
| Messages related to system startup are logged here.                                                                                                                                         | <code>/var/log/boot.log</code> |

## □ 10.3. Reviewing Syslog Files

### □ Objectives

After completing this section, students should be able to interpret entries in relevant syslog files to troubleshoot problems or review system status.

## ▮ Syslog files

### Reviewing Syslog Files

按照消息的类型和优先级分类

Many programs use the *syslog* protocol to log events to the system. Each log message is categorized by a facility (the type of message) and a priority (the severity of the message). The facilities which are available are documented by the `rsyslog.conf(5)` man page.

The eight priorities are also standardized and ranked as follows:

▮

### Overview of syslog priorities

| Code | Priority | Severity                          |
|------|----------|-----------------------------------|
| 0    | emerg    | System is unusable.               |
| 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 `rsyslogd` service uses the facility and priority of log messages to determine how to handle them.

This is configured by the file `/etc/rsyslog.conf` and by `*.conf` files

in `/etc/rsyslog.d`. Programs and administrators can change `rsyslogd` configuration in a way that will not be overwritten by updates to `rsyslog` by putting customized files with a `.conf` suffix in the `/etc/rsyslog.d` directory.

The `#### RULES ####` section of `/etc/rsyslog.conf` contains directives that define where log messages are saved. The left side of each line indicates the facility and severity of the log message the directive matches. The `rsyslog.conf` file can contain the character `*` as a wild card in the facility and severity field, where it either stands for all facilities or all severities. The right side of each line indicates what file to save the log message in. Log messages are normally saved in files in the `/var/log` directory.

允许使用通配符

### Note

Log files are maintained by the `rsyslog` service, and the `/var/log` directory contains a variety of log files specific to certain services. For example, the Apache Web Server or Samba write their own log files into a corresponding subdirectory of the `/var/log` directory.

A message handled by `rsyslog` can appear in multiple different log files. To prevent that, the severity field can be set to `none`, which means that none of the messages directed to this facility are added to the specified log file.

Instead of logging syslog messages to a file, they can be printed to the terminals of all logged-in users. In the default `rsyslog.conf` file, this is done for all messages that have "emerg" priority.



## Sample rules section of `rsyslog.conf`

```
RULES

Log all kernel messages to the console.
Logging much else clutters up the screen.
#kern.* /dev/console

Log anything (except mail) of level info or higher.
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 **`rsyslog.conf`** file is documented by the **`rsyslog.conf(5)`** man page and by extensive HTML documentation in **`/usr/share/doc/rsyslog-*/manual.html`** contained in the *rsyslog-doc*, which is available from the Red Hat Enterprise Linux 7 software channel, but not included on the installation medium.

## Log file rotation

Logs are "rotated" by the **logrotate** utility to keep them from filling up the file system containing **/var/log/**. When a log file is rotated, it is renamed with an extension indicating the date on which it was rotated: the old **/var/log/messages** file may become **/var/log/messages-20141030** if it is rotated on October 30, 2014. Once the old log file is rotated, a new log file is created and the service that writes to it is notified. After a certain number of rotations, typically after four weeks, the old log file is discarded to free disk space. **A cron job runs the logrotate program daily to see if any logs need to be rotated.** Most log files are rotated **weekly**, but logrotate rotates some faster, or slower, or when they reach a certain size. Configuration of logrotate is not covered in this course. For more information, see the **logrotate(8)** man page.

## Analyze a syslog entry

最新消息是在日志文件底部

**The system logs written by rsyslog start with the oldest message on top and the newest message at the end of the log file.** All log entries in log files managed by **rsyslog** are recorded in a standard format. The following example will explain the anatomy of a log file message in the **/var/log/secure** log file: **标准的格式记录**

```
1 Feb 11 20:11:48 2 localhost 3 sshd[1433]: 4 Failed password for student from 172.25.0.10 port 59344 ssh2
```

- 1 The time stamp when the log entry was recorded.
- 2 The host from which the log message was sent.
- 3 The program or process that sent the log message.
- 4 The actual message sent.

## Monitor a log file with tail

监视日志文件的变化

It is especially helpful for reproducing problems and issues to monitor one or more log files for events. The **tail -f /path/to/file** command outputs the last 10 lines of the file specified and continues to output new lines as they get written to the monitored file.

To monitor for failed login attempts on one terminal, run **ssh** as user **root** while a user tries to log in to the **serverX** machine:

```
[root@serverX ~]$ tail -f /var/log/secure
```

```
...
```

```
Feb 10 09:01:13 localhost sshd[2712]: Accepted password for root from
172.25.254.254 port 56801 ssh2

Feb 10 09:01:13 localhost sshd[2712]: pam_unix(sshd:session): session opened for
user root by (uid=0)
```

## Send a syslog message with `logger` 使用logger发送日志消息

The `logger` command can send messages to the `rsyslog` service. By default, it sends the message to the facility user with severity notice (`user.notice`) unless specified otherwise with the `-p` option. It is especially useful to test changes to the `rsyslog` configuration.

To send a message to `rsyslogd` that gets recorded in the `/var/log/boot.log` log file, execute: 测试rsyslog文件的变化

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

### References

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

`rsyslog` Manual

- o `/usr/share/doc/rsyslog-*/manual.html` provided by the `rsyslog-doc` package

Additional information may be available in the *Red Hat Enterprise Linux System Administrator's Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

## 10.4. Practice: Finding Log Entries 演示

### Guided exercise

In this lab, you will reconfigure `rsyslog` to write specific messages to a new log file.

Outcomes:

The `rsyslog` service writes all messages with priority debug to

the `/var/log/messages-debug` log file for temporary troubleshooting purposes.

1. Configure `rsyslog` on serverX to log all messages with severity debug in the newly created log file `/var/log/messages-debug` by adding the `rsyslog` configuration file `/etc/rsyslog.d/debug.conf`. Verify that a generated debug log message with the `logger` command arrives in the `/var/log/messages-debug` log file.
  - a. Change the `rsyslog` configuration to log all messages with severity debug to `/var/log/messages-debug` on serverX by adding the `/etc/rsyslog.d/debug.conf` file.

```
[root@serverX ~]# echo "*debug
/var/log/messages-debug" >/etc/rsyslog.d/debug.
conf
```

- b. Restart the `rsyslog` service on serverX.

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

2. Generate a debug log message with the **logger** command and verify that the message gets logged to the log file **/var/log/messages-debug** with the **tail** command on serverX.

- a. Monitor the **/var/log/messages-debug** with the **tail** command on serverX.

```
[root@serverX ~]# tail -f /var/log/messages-debug
```

- b. On a separate terminal window, use the **logger** command to generate a debug message on serverX.

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

- c. Switch back to the terminal still running the **tail -f /var/log/messages-debug** command and verify the message sent with the **logger** command shows up.

```
d. [root@serverX ~]# tail -f /var/log/messages-debug
```

```
e. ...
```

```
Feb 13 10:37:44 localhost root: Debug Message Test
```

## 10.5. Reviewing systemd Journal Entries

### Objectives

After completing this section, students should be able to find and interpret log entries in the systemd journal to troubleshoot problems or review system status.

### Finding events with `journalctl`

#### Reviewing systemd Journal Entries

The systemd journal stores logging data in **a structured, indexed binary file**. This data includes extra information about the log event. For syslog events, this can include the facility and priority of the original message, for example.

#### Important

In Red Hat Enterprise Linux 7, the systemd journal is stored in **/run/log** by default, and **its contents are cleared after a reboot**. This setting can be changed by the system administrator and is discussed elsewhere in this course.

The **journalctl** command shows the full system journal, starting with the oldest log entry, when run as root user:

```
[root@serverX ~]# journalctl
```

```
Feb 13 10:01:01 server1 run-parts(/etc/cron.hourly) [8678]: starting
Oyum-hourly.cron
```

```
Feb 13 10:01:01 server1 run-parts(/etc/cron.hourly) [8682]: finished
Oyum-hourly.cron

Feb 13 10:10:01 server1 systemd[1]: Starting Session 725 of user root.

Feb 13 10:10:01 server1 systemd[1]: Started Session 725 of user root.

Feb 13 10:10:01 server1 CROND[8687]: (root) CMD (/usr/lib64/sa/sa1 1 1)
```

The **journalctl** command highlights in **bold text messages of priority notice or warning**, and messages of priority **error and higher are highlighted in red**. 黑体-警告, 红色字体-错误

**The key** to successfully using the journal for troubleshooting and auditing is to limit the journal searches to only show relevant output. In the following paragraphs, various different strategies to reduce the output of journal queries will be introduced.

By default, **journalctl -n** shows the last 10 log entries. It takes an optional parameter for how many of the last log entries should be displayed. To display the last 5 log entries, run:

```
[root@serverX ~]# journalctl -n 5
```

When troubleshooting problems, it is useful to **filter** the output of the journal by **priority** of the journal entries. The **journalctl -p** takes either **the name or the number** of the known priority levels and shows the given levels and all higher-level entries. The priority levels known 按优先级过滤 to **journalctl** are debug, info, notice, warning, err, crit, alert, and emerg.

To filter the output of the **journalctl** command to only list any log entry of priority err or above, run:

```
[root@serverX ~]# journalctl -p err
```

Similar to the **tail -f** command, **journalctl -f** outputs the last 10 lines of the journal and continues to output new journal entries as they get written to the journal.

```
[root@serverX ~]# journalctl -f
```

When looking for specific events, it is useful to limit the output to **a specific time frame**.

The **journalctl** command has two options to limit the output to a specific time range, the **--since** and **--until** options. Both options take a time parameter in the 时间范围 format **YYYY-MM-DD hh:mm:ss**. If the date is omitted, the command assumes the date is today, and if the time part is omitted, the whole day starting at 00:00:00 is assumed. Both options take **yesterday, today, and tomorrow** as valid parameters in addition to the date and time field.

Output all journal entries that got recorded today:

```
[root@serverX ~]# journalctl --since today
```

Output the journal entries from 10th February 2014 20:30:00 to 13th February 2014 12:00:00:

```
[root@serverX ~]# journalctl --since "2014-02-10 20:30:00"
--until "2014-02-13 12:00:00"
```

In addition to the **visible content of the journal**, there are fields attached to the log entries that can only be seen when verbose output is turned on. All of the displayed extra fields can be used to filter the

output of a journal query. This is useful to reduce the output of complex searches for certain events in the journal.

```
[root@serverX ~]# journalctl -o verbose
Thu 2014-02-13 02:06:00.409345 EST [s=0b47abff995149c191a8e539e18c3f9c;
i=d28;b=1ea26e84667848af9a4a2904a76ff9a5;m=4d6878ff5a;t=4f244525daa67;
x=880bc65783036719]

 _PRIORITY=6
 _UID=0
 _GID=0
 _BOOT_ID=1ea26e84667848af9a4a2904a76ff9a5
 _MACHINE_ID=4513ad59a3b442ffa4b7ea88343fa55f
 _CAP_EFFECTIVE=0000001fffffffffff
 _TRANSPORT=syslog
 SYSLOG_FACILITY=10
 SYSLOG_IDENTIFIER=sshd
 _COMM=sshd
 _EXE=/usr/sbin/sshd
 _SYSTEMD_CGROUP=/system.slice/sshd.service
 _SYSTEMD_UNIT=sshd.service
 _SELINUX_CONTEXT=system_u:system_r:sshd_t:s0-s0:c0.c1023
 _HOSTNAME=serverX
 _CMDLINE=sshd: root [priv]
 SYSLOG_PID=6833
 _PID=6833
 MESSAGE=Failed password for root from 172.25.X.10 port 59371 ssh2
 _SOURCE_REALTIME_TIMESTAMP=1392275160409345
```

Among the more useful options to search for lines relevant to a particular process or event are:

- **COMM** The name of the command
- **EXE** The path to the executable for the process
- **PID** The PID of the process
- **UID** The UID of the user running the process
- **SYSTEMD\_UNIT** The systemd unit that started the process

More than one of these can be combined. For example, the following query shows all journal entries related to processes started by the systemd unit file `sshd.service`, which also have PID 1182:

```
[root@serverX ~]# journalctl _SYSTEMD_UNIT=sshd.service
_PID=1182
```

#### Note

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

#### References

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

Additional information may be available in the *Red Hat Enterprise Linux System Administrator's Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

## 10.6. Practice: Finding Events With `journalctl`

### Guided exercise

In this lab, you will filter the `systemd` journal for specific criteria.

#### Outcomes:

Students will practice displaying the `systemd` journal output matching different criteria.

1. Output only `systemd` journal messages that originate from the `systemd` process that always runs with process id 1 on serverX.

```
[root@serverX ~]# journalctl _PID=1
```

2. Display all `systemd` journal messages that originate from a system service started with user id 81 on serverX.

```
[root@serverX ~]# journalctl _UID=81
```

3. Output the journal messages with priority `warning` and above on serverX.

```
[root@serverX ~]# journalctl -p warning
```

4. Create a `journalctl` query to show all log events recorded in the previous 10 minutes on serverX. The command assumes a current time of 9:15:00.

```
[root@serverX ~]# journalctl --since 9:05:00 --until
9:15:00
```

5. Display only the events originating from the `sshd` service with the system unit file `sshd.service` recorded since 9:00:00 this morning on serverX.

```
[root@serverX ~]# journalctl --since 9:00:00
_SYSTEMD_UNIT="sshd.service"
```

## 10.7. Preserving the `systemd` Journal 保留系统日志

### Objectives

After completing this section, students should be able to configure `systemd-journald` to store its journal on disk rather than in memory.

## Store the system journal permanently

By default, the systemd journal is kept in `/run/log/journal`, which means it is cleared when the system reboots. **The journal is a new mechanism in Red Hat Enterprise Linux 7**, and for most installations, a detailed journal that starts with the last boot is sufficient.

If the directory `/var/log/journal` exists, the journal will log to that directory instead. **The advantage of this is the historic data will be available immediately at boot**. However, even with a persistent journal, not all data will be kept forever. **The journal has a built-in log rotation mechanism that will trigger monthly**. In addition, by default, **the journal will not be allowed to get larger than 10% of the file system it is on, or leave less than 15% of the file system free**. These values can be tuned in `/etc/systemd/journald.conf`, and the current limits on the size of the journal are logged when the **systemd-journald process starts**, as can be seen by the following command, which shows the top two lines of `journalctl` output:

```
[root@serverX ~]# journalctl | head -2
-- Logs begin at Wed 2014-03-05 15:13:37 CST, end at Thu 2014-03-06 21:57:54 CST.
--
Mar 05 15:13:37 serverX.example.com systemd-journal[94]: Runtime journal is using
8.0M (max 277.8M, leaving 416.7M of free 2.7G, current limit 277.8M).
```

The systemd journal can be made persistent by creating the directory `/var/log/journal` as user root:

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

Ensure that the `/var/log/journal` directory is owned by the root user and group `systemd-journal`, and has the permissions `2755`.

```
[root@serverX ~]# chown root:systemd-journal
/var/log/journal
[root@serverX ~]# chmod 2755 /var/log/journal
```

Either a reboot of the system or sending the special signal `USR1` as user root to the `systemd-journald` process is required.

```
[root@serverX ~]# killall -USR1 systemd-journald
```

Since the systemd journal is now persistent across reboots, **journalctl -b** can reduce the output by **only showing the log messages since the last boot of the system**.

```
[root@serverX ~]# journalctl -b
```

← 仅显示上自上次启动时的日志信息

### Note

When debugging a system crash with a persistent journal, it is usually required to limit the journal query to the reboot before the crash happened. The `-b` option can be accompanied by a negative number indicating to how many prior system boots the output should be limited. For example, the `journalctl -b -1` limits the output to the previous boot.

### References

加速启动

日志轮转, 按月



`mkdir(1)`, `systemd-journald(1)`, and `killall(1)` man pages

Additional information may be available in the *Red Hat Enterprise Linux System Administrator's Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

## 10.8. Practice: Configure a Persistent systemd Journal

### Guided exercise

In this lab, students will make the systemd journal persistent.

#### Outcomes:

The **systemd** journal is written to disk.

1. Configure the systemd journal to be persistent across reboots.
  - a. Configure the directory `/var/log/journal` on serverX.

```
b. [root@serverX ~]# mkdir /var/log/journal
c. [root@serverX ~]# chown root:systemd-journal
 /var/log/journal
d. [root@serverX ~]# chmod 2755 /var/log/journal
```

- e. Send the **USR1** signal to the `systemd-journald` or reboot serverX.

```
[root@serverX ~]# killall -USR1 systemd-journald
```

USR1亦通常被用来告知应用程序重载配置文件

2. To verify the systemd journal is persistent, look for a new directory with the systemd journal log files that have been written to `/var/log/journal`. (The exact files which appear may vary on your system, but the directory should have similar contents to the following example.)

```
3. [root@serverX ~]# ls
 /var/log/journal/4513ad59a3b442ffa4b7ea88343fa55f
```

```
system.journal user-1000.journal
```

## 10.9. Maintaining Accurate Time

### Objectives

After completing this section, students should be able to maintain accurate time synchronization and time zone configuration to ensure correct timestamps in system logs.

### Set local clocks and time zone

设置本地时钟和时区

#### Maintaining accurate time

Correct synchronized system time is very important for log file analysis across multiple systems.

The **Network Time Protocol (NTP)** is a standard way for machines to provide and obtain correct time information **on the Internet**. A machine may get accurate time information from **public NTP services** on the Internet such as the NTP Pool Project. A **high-quality** hardware clock to serve accurate time to local clients is another option.

高品质的

正确的系统时间对于系统及分析日志非常重要

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

```
[student@serverX ~]$ timedatectl

Local time: Thu 2014-02-13 02:16:15 EST
Universal time: Thu 2014-02-13 07:16:15 UTC
RTC time: Thu 2014-02-13 07:16:15
Timezone: America/New_York (EST, -0500)
NTP enabled: yes
NTP synchronized: no
RTC in local TZ: no
DST active: no
Last DST change: DST ended at
 Sun 2013-11-03 01:59:59 EDT
 Sun 2013-11-03 01:00:00 EST
Next DST change: DST begins (the clock jumps one hour forward) at
 Sun 2014-03-09 01:59:59 EST
 Sun 2014-03-09 03:00:00 EDT
```

A database with known time zones is available and can be listed with:

```
[student@serverX ~]$ timedatectl list-timezones
Africa/Abidjan
Africa/Accra
Africa/Addis_Ababa
Africa/Algiers
Africa/Asmara
Africa/Bamako
...
```

**Time zone names are based on the public "tz" (or "zoneinfo") time zone database maintained by IANA.**

Time zones are named based on continent or ocean, then typically but not always the largest city within the time zone region. For example, most of the US Mountain time zone is "America/Denver." Selecting the correct name can be non-intuitive in cases where localities inside the time zone have different daylight saving time rules. For example, in the USA, much of the state of Arizona (US Mountain time) does not have a daylight saving time adjustment at all and is in the time zone "America/Phoenix."

The command **tzselect** is useful for identifying correct zoneinfo time zone names. It interactively prompts the user with questions about the system's location, and outputs the name of the correct time zone. It does not make any changes to the time zone setting of the system.

The system setting for the current time zone can be adjusted as user root:

```
[root@serverX ~]# timedatectl set-timezone America/Phoenix
[root@serverX ~]# timedatectl

Local time: Thu 2014-02-13 00:23:54 MST
Universal time: Thu 2014-02-13 07:23:54 UTC
RTC time: Thu 2014-02-13 07:23:53
Timezone: America/Phoenix (MST, -0700)
NTP enabled: yes
NTP synchronized: no
RTC in local TZ: no
DST active: n/a
```

也可用于更改当前系统时间

To **change the current time and date settings with the `timedatectl` command**, the **`set-time`** option is available. The time is specified in the **"YYYY-MM-DD hh:mm:ss" format**, where either date or time can be **omitted**. To change the time to 09:00:00, run:

```
[root@serverX ~]$ timedatectl set-time 9:00:00
[root@serverX ~]$ timedatectl

Local time: Thu 2014-02-13 09:00:27 MST
Universal time: Thu 2014-02-13 16:00:27 UTC
RTC time: Thu 2014-02-13 16:00:28
Timezone: America/Phoenix (MST, -0700)
NTP enabled: yes
NTP synchronized: no
RTC in local TZ: no
DST active: n/a
```

The **`set-ntp`** option enables or disables NTP synchronization for automatic time adjustment. The option requires either a **`true`** or **`false`** argument to turn it on or off. To turn on NTP synchronization, run:

```
[student@desktopX ~]$ timedatectl set-ntp true
```

## Configuring and monitoring chronyd

The **`chronyd`** service keeps the usually-inaccurate local **hardware clock (RTC)** on track by synchronizing it to the configured NTP servers, or if no network connectivity is available, to the calculated RTC clock **drift**, which is recorded in the **`driftfile`** specified in the **`/etc/chrony.conf`** configuration file.

By default, **`chronyd`** uses servers from the NTP Pool Project for the time synchronization and does not need additional configuration. **It may be useful to change the NTP servers when the machine in question is on an isolated network.**

在一个独立的网络中更改NTP服务器是非常有用的

The quality of an NTP time source is determined by the **stratum value** reported by the time source. The **stratum** determines the number of hops the machine is away from a high-performance reference clock. The reference clock is a **stratum 0** time source. An NTP server directly attached to it is a **stratum 1**, while a machine synchronizing time from the NTP server is a **stratum 2** time source.

There are two categories of time sources that can be configured in the **/etc/chrony.conf** configuration file, **server** and **peer**. The **server** is one stratum above the local NTP server, and the **peer** is at the same stratum level. More than one **server** and more than one **peer** can be specified, 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, a series of options for the server can be listed. It is recommended to use the **iburst** option, because after the service starts, four measurements are taken in a short time period for a more accurate initial clock synchronization.

To reconfigure the **chronyd** server to synchronize with classroom.example.com instead of the default servers configured in the **/etc/chrony.conf**, remove the other server entries and replace them with the following configuration file entry:

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

After pointing **chronyd** to the local time source, classroom.example.com, the service needs to be restarted:

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

The **chronyc** command acts as a client to the **chronyd** service. After setting up NTP synchronization, it is useful to verify the NTP server was used to synchronize the system clock. This can be achieved with the **chronyc sources** command or, for more verbose output with additional explanations about the output, **chronyc sources -v**:

```
[root@serverX ~]$ chronyc sources -v
210 Number of sources = 1

.-- Source mode '^' = server, '=' = peer, '#' = local clock.
/ .- Source state '*' = current synced, '+' = combined , '-' = not combined,
| / '?' = unreachable, 'x' = time may be in error, '~' = time too variable.
||
|| .- xxxx [yyyy] +/- zzzz
|| / 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 8 6 17 23 -497ns[-7000ns] +/- 956us
```

The **\*** character in the **S** (Source state) field indicates that the classroom.example.com server has been used as a time source and is the NTP server the machine is currently synchronized to.

#### Note

Red Hat Enterprise Linux 6 and earlier use **ntpd** and **ntpq** to manage the NTP configuration. Further information may be found in the documentation for Red Hat Enterprise Linux 6.

#### References

**timedatectl**(1), **tzselect**(8), **chronyd**(8), **chrony.conf**(5), and **chronyc**(1) man pages

Additional information may be available in the *Red Hat Enterprise Linux System Administrator's Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

[NTP Pool Project](#)

[Time Zone Database](#)

## 10.10. Practice: Adjusting System Time

### Quiz

The steps to set up the correct time zone and adjust system clocks by using **timedatectl** and **chronyd** follow. Indicate the order in which the steps should be taken.

- Review current settings with **timedatectl**.
- Identify the appropriate time zone with the **timedatectl list-timezones** command.
- Set the correct time zone with **timedatectl set-timezone**.
- Turn on NTP synchronization.**
- Point **chronyd** to a new time source by adjusting **/etc/chrony.conf**.
- Restart the **chronyd** service.
- Verify the clock was synchronized against the newly specified NTP source with the **chronyc sources** command.

## 10.11. Lab: Analyzing and Storing Logs

演示给同学们

### Performance checklist

In this lab, students will change the time zone and log all authentication failure log entries into a separate file.

#### Outcomes:

The time zone setting on the serverX machine is adjusted; all **systemd** journal entries recorded in a given time frame are displayed; all syslog messages with the authpriv facility and severity alert are logged into a separate log file.

#### Before you begin...

Reset your serverX system.

1. Your serverX machine has been relocated to Jamaica. Change the time zone on the serverX machine to Jamaica and verify the time zone has been changed properly.

- a. Identify the correct time zone for Jamaica on serverX.

```
b. [root@serverX ~]# timedatectl list-timezones
c. Africa/Abidjan
d. Africa/Accra
e. Africa/Addis_Ababa
f. Africa/Algiers
g. Africa/Asmara
h. ...
i. America/Jamaica
j. ...
```

- k. Change the time zone to Jamaica on serverX.

```
[root@serverX ~]# timedatectl set-timezone
America/Jamaica
```

- l. Verify the time zone has been properly set on serverX.

```
m. [root@serverX ~]# timedatectl
n. Local time: Thu 2014-02-13 11:16:59 EST
o. Universal time: Thu 2014-02-13 16:16:59 UTC
p. RTC time: Thu 2014-02-13 16:17:00
q. Timezone: America/Jamaica (EST, -0500)
r. NTP enabled: yes
s. NTP synchronized: no
t. RTC in local TZ: no
```

```
DST active: n/a
```

2. Display all **systemd** journal entries recorded in the last 30 minutes on serverX.

Assuming the current time is 9:30:00, the following command would be used

```
[root@serverX ~]# journalctl --since 9:00:00 --until
9:30:00
```

3. Configure **rsyslogd** by adding a rule to the newly created configuration file **/etc/rsyslog.d/auth-errors.conf** to log all security and authentication messages with the priority alert and higher to the **/var/log/auth-errors** file as well. Test the newly added log directive with the **logger** command.

- a. Add the directive to log **authpriv.alert** syslog messages to the **/var/log/auth-errors** file in the **/etc/rsyslog.d/auth-errors.conf** configuration file.

```
[root@serverX ~]# echo "authpriv.alert
/var/log/auth-errors" >/etc/rsyslog.d/auth-err
ors.conf
```

- b. Restart the **rsyslog** service on serverX.

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

- c. Monitor the newly created log file **/var/log/auth-errors** on serverX for changes in a different terminal window.

```
[root@serverX ~]# tail -f /var/log/auth-errors
```

- d. Use the **logger** to create a new log entry to the **/var/log/auth-errors** on serverX.

```
[root@serverX ~]# logger -p authpriv.alert
"Logging test authpriv.alert"
```

- e. Verify the message sent to syslog with the **logger** command appears in the **/var/log/auth-errors** on serverX in the terminal running **tail -f /var/log/auth-errors**.

```
f. [root@serverX ~]# tail -f /var/log/auth-errors
```

```
Feb 13 11:21:53 server1 root: Logging test authpriv.alert
```

# □ Chapter 11. Managing Red Hat Enterprise Linux Networking

11.1. Networking Concepts

11.2. Practice: Networking Concepts

11.3. Validating Network Configuration

11.4. Practice: Examining Network Configuration

11.5. Configuring Networking with `nmcli`

11.6. Practice: Configuring Networking with `nmcli`

11.7. Editing Network Configuration Files

11.8. Practice: Editing Network Configuration Files

11.9. Configuring Host Names and Name Resolution

11.10. Practice: Configuring Host Names and Name Resolution

11.11. Lab: Managing Red Hat Enterprise Linux Networking

## Introduction

| Overview   |                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Goal       | To <b>configure basic IPv4 networking</b> on Red Hat Enterprise Linux systems.                                                                                                                                                                                                                                                                                                                                           |
| Objectives | <ul style="list-style-type: none"><li>○ Explain fundamental concepts of computer networking.</li><li>○ Test and review current network configuration with basic utilities.</li><li>○ Manage network settings and devices with <b>nmcli</b> and <b>NetworkManager</b>.</li><li>○ Modify network settings by editing the configuration files.</li><li>○ Configure and test system host name and name resolution.</li></ul> |
| Sections   | <ul style="list-style-type: none"><li>○ Networking Concepts (and Practice)</li><li>○ Validating Network Configuration (and Practice)</li><li>○ Configuring Networking with <b>nmcli</b> (and Practice)</li><li>○ Editing Network Configuration Files (and Practice)</li><li>○ Configuring Host Names and Name Resolution (and Practice)</li></ul>                                                                        |
| Lab        | <ul style="list-style-type: none"><li>○ Managing Red Hat Enterprise Linux Networking</li></ul>                                                                                                                                                                                                                                                                                                                           |

## 11.1. Networking Concepts

### Objectives

After completing this section, students should be able to explain fundamental concepts of computer networking.

### IPv4 networking

TCP/IP standards follow a four-layer network model specified in RFC1122.

- Application



Each application has specifications for communication so that clients and servers may communicate across platforms. Common protocols include SSH (remote login), HTTPS (secure web), NFS or CIFS (file sharing), and SMTP (electronic mail delivery).

- **Transport**

Transport protocols are TCP and UDP. TCP is a reliable connection-oriented communication, while UDP is a connectionless *datagram* protocol. Application protocols use TCP or UDP ports. A list of well-known and registered ports can be found in the `/etc/services` file.

When a packet is sent on the network, the combination of the service port and IP address forms a socket. Each packet has a source socket and a destination socket. This information can be used when monitoring and filtering.

- **Internet**

The Internet, or network layer, carries data from the source host to the destination host. Each host has an IP address and a prefix used to determine network addresses. Routers are used to connect networks.

ICMP is a control protocol at this layer. Instead of ports, it has types. The **ping** utility is an example of using ICMP packets to test connectivity. **ping** sends an ICMP ECHO\_REQUEST packet. A successful **ping** receives an ICMP ECHO\_REPLY acknowledgment. An unsuccessful **ping** may receive ICMP error messages such as "destination unreachable" or may not receive any response.

- **Link**

The link, or media access, layer provides the connection to physical media. The most common types of networks are wired Ethernet (802.3) and wireless WLAN (802.11). Each physical device has a hardware address (MAC) which is used to identify the destination of packets on the local network segment.

*Figure 11.1: IPv4 addresses and netmasks*

**IPv4 addresses**

An IPv4 address is a 32-bit number, normally expressed in decimal as four *octets* ranging in value from 0 to 255, separated by dots. The address is divided into two parts: the *network part* and the *host part*. All hosts on the same subnet, which can talk to each other directly without a router, have the same

network part; the network part identifies the subnet. No two hosts on the same subnet can have the same host part; the host part identifies a particular host on a subnet.

In the modern Internet, the size of an IPv4 subnet is variable. To know which part of an IPv4 address is the network part and which is the host part, an administrator must know the *netmask* which is assigned to the subnet. The netmask indicates how many bits of the IPv4 address belong to the subnet. The more bits that are available for the host part, the more hosts can be on the subnet.

The lowest possible address on a subnet (host part is all zeros in binary) is sometimes called the *network address*. The highest possible address on a subnet (host part is all ones in binary) is used for broadcast messages in IPv4, and is called the *broadcast address*.

Network masks are expressed in two forms. The older syntax for a netmask which uses 24 bits for the network part would read 255.255.255.0. A newer syntax, called CIDR notation, would specify a *network prefix* of /24. Both forms convey the same information; namely, how many leading bits in the IP address contribute to its network address.

The examples which follow illustrate how the IP address, prefix (netmask), network part, and host part are related.

▮

#### Calculating the network address for 192.168.1.107/24

|                |                     |                                            |
|----------------|---------------------|--------------------------------------------|
| Host addr      | 192.168.1.107       | <b>11000000.10101000.00000001.01101011</b> |
| Network prefix | /24 (255.255.255.0) | <b>11111111.11111111.11111111.00000000</b> |
| Network addr   | 192.168.1.0         | <b>11000000.10101000.00000001.00000000</b> |
| Broadcast addr | 192.168.1.255       | <b>11000000.10101000.00000001.11111111</b> |

▮

#### Calculating the network address for 10.1.1.18/8

|                |                |                                            |
|----------------|----------------|--------------------------------------------|
| Host addr      | 10.1.1.18      | <b>00001010.00000001.00000001.00010010</b> |
| Network prefix | /8 (255.0.0.0) | <b>11111111.00000000.00000000.00000000</b> |
| Network addr   | 10.0.0.0       | <b>00001010.00000000.00000000.00000000</b> |
| Broadcast addr | 10.255.255.255 | <b>00001010.11111111.11111111.11111111</b> |

▮

#### Calculating the network address for 172.16.181.23/19

|                |                     |                                            |
|----------------|---------------------|--------------------------------------------|
| Host addr      | 172.168.181.23      | <b>10101100.10101000.10110101.00010111</b> |
| Network prefix | /19 (255.255.224.0) | <b>11111111.11111111.11100000.00000000</b> |
| Network addr   | 172.168.160.0       | <b>10101100.10101000.10100000.00000000</b> |
| Broadcast addr | 172.168.191.255     | <b>10101100.10101000.10111111.11111111</b> |

The special address 127.0.0.1 always points to the local system ("localhost"), and the network 127.0.0.0/8 belongs to the local system, so that it can talk to itself using network protocols.

*Figure 11.2: Network routing and DNS concepts*

#### **IPv4 routing**

Whether using IPv4 or IPv6, network traffic needs to move from host to host and network to network. Each host has a *routing table*, which tells it how to route traffic for particular networks. The routing table entries will list a destination network, which interface to send the traffic out, and the IP address

of any intermediate router that is required to relay the message to its final destination. The routing table entry which matches the destination of the network traffic is used to route it. If two entries match, the one with the longest prefix is used.

If the network traffic does not match a more specific route, the routing table usually has an entry for a *default route* to the entire IPv4 Internet, 0.0.0.0/0. This default route points to a *router* on a reachable subnet (that is, on a subnet that has a more specific route in the host's routing table).

If a router receives traffic that is not addressed to it, instead of ignoring it like a normal host, it *forwards* the traffic based on its own routing table. This may send the traffic directly to the destination host (if the router happens to be on the destination's subnet), or it may be forwarded on to another router. This process of forwarding continues until the traffic reaches its final destination.

□

### Example routing table

| Destination                  | Interface     | Router (if needed) |
|------------------------------|---------------|--------------------|
| 192.0.2.0/24                 | <b>wlo1</b>   |                    |
| 192.168.5.0/24               | <b>enp3s0</b> |                    |
| 0.0.0.0/0 ( <i>default</i> ) | <b>enp3s0</b> | 192.168.5.254      |

In this example, traffic headed for the IP address 192.0.2.102 from this host will be transmitted directly to that destination via the **wlo1** wireless interface, because it matches the 192.0.2.0/24 route most closely. Traffic for the IP address 192.168.5.3 will be transmitted directly to that destination via the **enp3s0** Ethernet interface, because it matches the 192.168.5.0/24 route most closely.

Traffic to the IP address 10.2.24.1 will be transmitted out the **enp3s0** Ethernet interface to a router at 192.168.5.254, which will forward that traffic on to its final destination. That traffic matches the 0.0.0.0/0 route most closely, as there is not a more specific route in the routing table of this host. The router will use its own routing table to determine where to forward that traffic to next.

### Names and IP addresses

The IP protocol uses addresses to communicate, but human beings would rather work with names than long and hard-to-remember strings of numbers. *DNS*, the Domain Name System, is a distributed network of servers that maps host names to IP addresses. In order for name service to work, the host needs to be pointed at a *nameserver*. This nameserver does not need to be on the same subnet; it just needs to be reachable by the host.

### DHCP or static network configuration

Many systems are configured to obtain network settings automatically at boot time. The local configuration files indicate that DHCP should be used and a separate client service queries the network for a server and obtains a lease for network settings.

If a DHCP server is not available, the system must use a *static* configuration where the network settings are read from a local configuration file. The correct network settings are obtained from the network administrator or architecture team to ensure there are no conflicts with other systems.

Since DHCP uses the hardware address to track assignments, only one address may be assigned per interface with DHCP. Multiple static addresses may be assigned to a single interface. This practice is common in systems hosting services for multiple clients, such as HTTP IP-based hosting. Red Hat Enterprise Linux interfaces typically have an IPv4 address and an IPv6 local-link address, but may have more addresses assigned.

## Network interface names

枚举

Traditionally, network interfaces in Linux are enumerated as **eth0**, **eth1**, **eth2**, and so on. However, the mechanism which sets these names can cause changes to which interface gets which name as devices are added and removed. The default naming behavior in Red Hat Enterprise Linux 7 is to assign fixed names based on firmware, device topology, and device type. Interface names have the following characters:

当设备添加或移除时会引起接口名字的改变

RHEL7为接口指派了固定的名字，命名基于固件、设备的拓扑和设备的类型

- Ethernet interfaces begin with **en**, WLAN interfaces begin with **wl**, and WWAN interfaces begin with **ww**.
- The next character(s) represents the type of adapter with an **o** for on-board, **s** for hotplug slot, and **p** for PCI geographic location. Not used by default but also available to administrators, an **x** is used to incorporate a MAC address.
- Finally, a number **N** is used to represent an index, ID, or port.
- If the fixed name cannot be determined, the traditional names such as **ethN** will be used.

内置的

For example, the first embedded network interface may be named **eno1** and a PCI card network interface may be named **enp2s0**. The new names make it easier to distinguish the relationship between a port and its name if the user knows both, but the trade off is that users cannot assume a system with one interface calls that interface **eth0**.

新的命名非常容易区别端口和设备名字之间的关系，如果用户知道的话

### Note

Network interface naming can be overridden. If the administrator has installed and enabled the **biosdevname** package or set customized **udev** device naming rules, those settings will override the default naming scheme. Depending on support for **biosdevname** in the system BIOS, names such as **em1**, **em2**, etc. may be used for on-board network cards (corresponding to their names on the chassis). PCI(e) cards are represented with **pYpX** (e.g., **p4p1**), where **Y** is the PCI slot number and **X** is the number for the port on that specific card.

### References

**services(5)**, **ping(8)**, **biosdevname(1)**, and **udev(7)** man pages

Additional information may be available in the chapters on configuring networking and consistent network device naming in the *Red Hat Enterprise Linux Networking Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

## 11.2. Practice: Networking Concepts

### Quiz

Match the following items to their counterparts in the table.

| Network settings                                                               | Correctness          |
|--------------------------------------------------------------------------------|----------------------|
| IP address: 172.17.0.351/16<br>Gateway: 172.17.0.1<br>DNS server: 172.17.0.254 | Invalid IPv4 address |

| Network settings                                                                | Correctness                             |
|---------------------------------------------------------------------------------|-----------------------------------------|
| IP address: 10.1.2.3/24<br>Gateway: 10.1.2.1<br>DNS server: 172.17.4.53         | This configuration is feasible.         |
| IP address: 192.168.7.0/24<br>Gateway: 192.168.7.1<br>DNS server: 192.168.0.254 | IP address cannot be a network address. |
| IP address: 10.4.5.6/24<br>Gateway: 10.4.6.1<br>DNS server: 192.168.0.254       | Gateway is not on the same subnet.      |
| IP address: 172.17.23.5/16<br>Gateway: 172.17.0.1                               | Name resolution is not configured.      |

## 11.3. Validating Network Configuration

### Objectives

After completing this section, students should be able to test and review current network configuration with basic utilities.

### Displaying IP addresses

*Validating network configuration*

The `/sbin/ip` command is used to show device and address information.

```
[student@desktopX ~]$ ip addr show eth0

2: eth0: <BROADCAST,MULTICAST, 1 UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state
UP qlen 1000

2 link/ether 52:54:00:00:00:0a brd ff:ff:ff:ff:ff:ff

3 inet 172.25.0.10/24 brd 4 172.25.0.255 scope global eth0
valid_lft forever preferred_lft forever

5 inet6 fe80::5054:ff:fe00:b/64 scope link
valid_lft forever preferred_lft forever
```

- 1 An active interface has the status of **UP**.
- 2 The link line specifies the hardware (MAC) address of the device.
- 3 The inet line shows the IPv4 address and prefix.
- 4 The broadcast address, scope, and device name are also on this line.
- 5 The inet6 line shows IPv6 information.

The **ip** command may also be used to **show statistics about network performance**. The received (RX) and transmitted (TX) packets, errors, and dropped counters can be used to identify network issues caused by congestion, low memory, and overruns.

```
[student@desktopX ~]$ ip -s link show eth0
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
qlen 1000
link/ether 52:54:00:00:00:0a brd ff:ff:ff:ff:ff:ff
 RX: bytes packets errors dropped overrun mcast
 269850 2931 0 0 0 0
 TX: bytes packets errors dropped carrier collsns
 300556 3250 0 0 0 0
```

## ▮ Troubleshooting routing

The **/sbin/ip** command is also used to show routing information.

```
[student@desktopX ~]$ ip route
default via 172.25.0.254 dev eth0 proto static metric 1024
172.25.X.0/24 dev eth0 proto kernel scope link src 172.25.X.10
10.0.0.0/8 dev eth1 proto kernel scope link src 10.0.0.11
```

All packets destined for the 10.0.0.0/8 network will be sent directly to the destination through the device eth1. All packets destined for the 172.25.x.0/24 network will be sent directly to the destination through the device eth0. All other packets will be sent to the default router located at 172.25.x.254, and also through device eth0.

The **ping** command is used to test connectivity. The command will continue to run until a **Control-C** is pressed unless options are given to limit the number of packets sent.



```
[student@desktopX ~]$ ping -c3 172.25.X.254
```

To trace the path to a remote host, use either **traceroute** or **tracepath**. Both commands can be used to trace a path with **UDP packets**; however, many networks block UDP and ICMP traffic. The **traceroute** command has options to trace the path with **UDP (default), ICMP (-I), or TCP (-T)** packets, but may not be installed by default.

```
[student@desktopX ~]$ tracepath access.redhat.com
...
4: 71-32-28-145.rcmt.qwest.net 48.853ms asymm 5
5: dcp-brdr-04.inet.qwest.net 100.732ms asymm 7
6: 206.111.0.153.ptr.us.xo.net 96.245ms asymm 7
7: 207.88.14.162.ptr.us.xo.net 85.270ms asymm 8
8: ae1d0.cir1.atlanta6-ga.us.xo.net 64.160ms asymm 7
9: 216.156.108.98.ptr.us.xo.net 108.652ms
10: bu-ether13.atlmgamq46w-bcr00.tbone.rr.com 107.286ms asymm 12
...
```

Each line in the output of **tracepath** represents a router or *hop* that the packet passes through between the source and the final destination. Additional information is provided as available, including the round trip timing (RTT) and any changes in the maximum transmission unit (MTU) size.

## ▣ Troubleshooting ports and services

TCP services use sockets as end points for communication and are made up of an IP address, protocol, and port number. Services typically listen on standard ports while clients use a random available port. Well-known names for standard ports are listed in the **/etc/services** file.

The **ss** command is used to display socket statistics. It is similar to the **netstat** command, which is also available but may not be installed by default.

```
[student@desktopX ~]$ ss -ta
```

| State  | Recv-Q | Send-Q | Local Address:Port       | Peer Address:Port    |
|--------|--------|--------|--------------------------|----------------------|
| LISTEN | 0      | 128    | *:sunrpc                 | *:*                  |
| LISTEN | 0      | 128    | <b>1</b> *:ssh           | *:*                  |
| LISTEN | 0      | 100    | <b>2</b> 127.0.0.1:smtp  | *:*                  |
| LISTEN | 0      | 128    | *:36889                  | *:*                  |
| ESTAB  | 0      | 0      | <b>3</b> 172.25.X.10:ssh | 172.25.254.254:59392 |
| LISTEN | 0      | 128    | :::sunrpc                | :::*                 |
| LISTEN | 0      | 128    | <b>4</b> :::ssh          | :::*                 |

```

LISTEN 0 100 5 ::::smtp :::*
LISTEN 0 128 :::34946 :::*

```

- 1 The port used for SSH is listening on all IPv4 addresses. The "\*" is used to represent "all" when referencing IPv4 addresses or ports.
- 2 The port used for SMTP is listening on the 127.0.0.1 IPv4 loopback interface.
- 3 The established SSH connection is on the 172.25.x.10 interface and originates from a system with an address of 172.25.254.254.
- 4 The port used for SSH is listening on all IPv6 addresses. The ":::" syntax is used to represent all IPv6 interfaces.
- 5 The port used for SMTP is listening on the ::1 IPv6 loopback interface.

▮

#### Options for **ss** and **netstat**

| Option | Description                                             |
|--------|---------------------------------------------------------|
| -n     | Show numbers instead of names for interfaces and ports. |
| -t     | Show TCP sockets.                                       |
| -u     | Show UDP sockets.                                       |
| -l     | Show only listening sockets.                            |
| -a     | Show all (listening and established) sockets.           |
| -p     | Show the process using the sockets.                     |

#### References

**ip-link(8)**, **ip-address(8)**, **ip-route(8)**, **ip(8)**, **ping(8)**, **tracepath(8)**, **traceroute(8)**, **ss(8)**, and **netstat(8)** man pages

Additional information may be available in the chapter on configuring networking in the *Red Hat Enterprise Linux Networking Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

## ▮ 11.4. Practice: Examining Network Configuration

### Guided exercise

In this lab, you will examine the network configuration of the current system.

#### Outcomes:

Identify the current network interfaces and basic network addresses.

## Before you begin...

Reset your serverX system.

1. Display the current IP address and netmask for all interfaces.

```
2. [student@serverX ~]$ ip addr
3. 1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
4. link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
5. inet 127.0.0.1/8 scope host lo
6. valid_lft forever preferred_lft forever
7. inet6 ::1/128 scope host
8. valid_lft forever preferred_lft forever
9. 2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
 qlen 1000
10. link/ether 52:54:00:00:00:0b brd ff:ff:ff:ff:ff:ff
11. inet 172.25.X.11/24 brd 172.25.X.255 scope global dynamic eth0
12. valid_lft 12704sec preferred_lft 12704sec
13. inet6 fe80::5054:ff:fe00:b/64 scope link
14. valid_lft forever preferred_lft forever
```

15. Display the statistics for the eth0 interface.

```
16. [student@serverX ~]$ ip -s link show eth0
17. 2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
 mode DEFAULT qlen 1000
18. link/ether 52:54:00:00:00:0b brd ff:ff:ff:ff:ff:ff
19. RX: bytes packets errors dropped overrun mcast
20. 418398 4588 0 0 0 0
21. TX: bytes packets errors dropped carrier collsns
22. 360733 1730 0 0 0 0
```

23. Display the routing information.

```
24. [student@serverX ~]$ ip route
25. default via 172.25.0.1 dev eth0 proto static metric 1024
26. 172.25.0.0/16 dev eth0 proto kernel scope link src 172.25.X.11
```

27. Verify that the router is accessible.

```
28. [student@serverX ~]$ ping -c3 172.25.0.1
29. PING 172.25.0.1 (172.25.0.1) 56(84) bytes of data.
30. 64 bytes from 172.25.0.1: icmp_seq=1 ttl=64 time=0.489 ms
```

```

31. 64 bytes from 172.25.0.1: icmp_seq=2 ttl=64 time=0.510 ms
32. 64 bytes from 172.25.0.1: icmp_seq=3 ttl=64 time=0.458 ms
33.
34. --- 172.25.0.1 ping statistics ---
35. 3 packets transmitted, 3 received, 0% packet loss, time 1999ms
36. rtt min/avg/max/mdev = 0.458/0.485/0.510/0.033 ms

```

37. Show all the hops between the local system and classroom.example.com.

```

38. [student@serverX ~]$ tracepath classroom.example.com
39. 1: classroom.example.com 0.522ms !H
40. Resume: pmtu 65535

```

41. Display the listening TCP sockets on the local system.

```

42. [student@serverX ~]$ ss -lt
43. State Recv-Q Send-Q Local Address:Port Peer Address:Port
44. LISTEN 0 128 *:55630 *:*
45. LISTEN 0 128 *:sunrpc *:*
46. LISTEN 0 128 *:ssh *:*
47. LISTEN 0 100 127.0.0.1:smtp *:*
48. LISTEN 0 128 :::sunrpc :::*
49. LISTEN 0 128 :::ssh :::*
50. LISTEN 0 128 :::33079 :::*

```

```
LISTEN 0 100 :::smtp :::*
```

## 11.5. Configuring Networking with `nmcli`

### Objectives

After completing this section, students should be able to manage network settings and devices with `nmcli` and NetworkManager.

### NetworkManager

NetworkManager is a daemon that monitors and manages network settings. In addition to the daemon, there is a GNOME Notification Area applet that provides network status information. Command-line and graphical tools talk to NetworkManager and save configuration files in the `/etc/sysconfig/network-scripts` directory.

A *device* is a network interface. A *connection* is a configuration used for a device which is made up of a collection of settings. Multiple connections may exist for a device, but only one may be active at a time. For example, a system may normally be connected to a network with settings provided by DHCP.

Occasionally, that system needs to be connected to a lab or data center network, which only uses static networking. Instead of changing the configuration manually, each configuration can be stored as a separate connection.

## ▣ Viewing network information with `nmcli`

*Configuring networking with nmcli*

To display a list of all connections, use `nmcli con show`. To list only the active connections, add the `--active` option.

```
[root@desktopX ~]# nmcli con show
```

| NAME        | UUID                                 | TYPE            | DEVICE |
|-------------|--------------------------------------|-----------------|--------|
| static-eth0 | f3e8dd32-3c9d-48f6-9066-551e5b6e612d | 802-3-ethernet  | eth0   |
| System eth0 | 5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03 | 802-3-ethernet  | --     |
| guest       | f601ca8a-6647-4188-a431-dab48cc63bf4 | 802-11-wireless | wlp3s0 |

```
[root@desktopX ~]# nmcli con show --active
```

| NAME        | UUID                                 | TYPE            | DEVICE |
|-------------|--------------------------------------|-----------------|--------|
| static-eth0 | f3e8dd32-3c9d-48f6-9066-551e5b6e612d | 802-3-ethernet  | eth0   |
| guest       | f601ca8a-6647-4188-a431-dab48cc63bf4 | 802-11-wireless | wlp3s0 |

Specify a connection ID (name) to see the details of that connection. The lowercase settings represent the configuration of the connection. Setting and property names are defined in the `nm-settings(5)` man page. The uppercase settings are active data.

```
[root@desktopX ~]# nmcli con show "static-eth0"
```

```
...
ipv4.method: manual
ipv4.dns: 172.25.254.254, 8.8.8.8
ipv4.dns-search:
ipv4.addresses: { ip = 172.25.X.10/24, gw = 172.25.X.254 }
ipv4.routes:
ipv4.ignore-auto-routes: no
ipv4.ignore-auto-dns: no
ipv4.dhcp-client-id: --
ipv4.dhcp-send-hostname: yes
ipv4.dhcp-hostname: --
ipv4.never-default: no
ipv4.may-fail: yes
ipv6.method: auto
...
```

The `nmcli` command can also be used to show device status and details.

```

[root@desktopX ~]# nmcli dev status
DEVICE TYPE STATE CONNECTION
eth0 ethernet connected static-eth0
wlp3s0 wifi connected guest
lo loopback unmanaged --

[root@desktopX ~]# nmcli dev show eth0
GENERAL.DEVICE: eth0
GENERAL.TYPE: ethernet
GENERAL.HWADDR: 52:54:00:00:00:0A
GENERAL.MTU: 1500
GENERAL.STATE: 100 (connected)
GENERAL.CONNECTION: static-eth0
GENERAL.CON-PATH: /org/freedesktop/NetworkManager/ActiveConnection/1
WIRED-PROPERTIES.CARRIER: on
IP4.ADDRESS[1]: ip = 172.25.X.10/24, gw = 172.25.X.254
IP4.DNS[1]: 172.25.254.254
IP6.ADDRESS[1]: ip = fe80::5054:ff:fe00:b/64, gw = ::

```

## Creating network connections with `nmcli`

When creating a new connection with `nmcli`, the order of the arguments is important. The common arguments appear first and must include the type and interface. Next, specify any type-specific arguments and finally specify the IP address, prefix, and gateway information. Multiple IP addresses may be specified for a single device. Additional settings such as a DNS server are set as modifications once the connection exists.

### Examples of creating new connections

Follow along with the next steps while your instructor discusses `nmcli` syntax.

1. Define a new connection named "default" which will autoconnect as an Ethernet connection on the eth0 device using DHCP.

```

2. [root@desktopX ~]# nmcli con add con-name "default"
 type ethernet ifname eth0

```

3. Create a new connection named "static" and specify the IP address and gateway. Do not autoconnect.

```

4. [root@desktopX ~]# nmcli con add con-name "static"
 ifname eth0 autoconnect no type ethernet ip4
 172.25.X.10/24 gw4 172.25.X.254

```

5. The system will autoconnect with the DHCP connection at boot. Change to the static connection.

```
6. [root@desktopX ~]# nmcli con up "static"
```

7. Change back to the DHCP connection.

```
8. [root@desktopX ~]# nmcli con up "default"
```

### Important

If the static connection is lost, the default connection will attempt to autoconnect. To administratively disable an interface and prevent any autoconnection, use **nmcli dev disconnect DEVICENAME**.

### Type options

Type options depend on the type used. An ethernet-type connection may optionally specify a MAC address for the connection. A wifi-type connection must specify the SSID and may specify additional options. Many other types are available, including bridge, bond, team, VPN, and VLAN. To view all the options, use **nmcli con add help**.

```
[root@desktopX ~]# nmcli con add help
Usage: nmcli connection add { ARGUMENTS | help }

ARGUMENTS := COMMON_OPTIONS TYPE_SPECIFIC_OPTIONS IP_OPTIONS

COMMON_OPTIONS:
 type <type>
 ifname <interface name> | "*"
 [con-name <connection name>
 [autoconnect yes|no]

 [save yes|no]

TYPE_SPECIFIC_OPTIONS:
 ethernet: [mac <MAC address>
 [cloned-mac <cloned MAC address>
 [mtu <MTU>

...

```

## Modifying network interfaces with `nmcli`

An existing connection may be modified with `nmcli con mod` arguments. The arguments are sets of key/value pairs. The key includes a setting name and a property name. Use `nmcli con show "<ID>"` to see a list of current values for a connection. The `nm-settings(5)` man page documents the setting and property names and usage.

```
[root@desktopX ~]# nmcli con show "static"
connection.id: static
connection.uuid: f3e8dd32-3c9d-48f6-9066-551e5b6e612d
connection.interface-name: eth0
connection.type: 802-3-ethernet
connection.autoconnect: yes
connection.timestamp: 1394905322
connection.read-only: no
...
```

### Examples of connection modifications

Follow along with the next steps while your instructor discusses `nmcli` syntax.

1. Turn off autoconnect.

```
2. [root@desktopX ~]# nmcli con mod "static"
 connection.autoconnect no
```

3. Specify a DNS server.

```
4. [root@desktopX ~]# nmcli con mod "static" ipv4.dns
 172.25.X.254
```

5. Some configuration arguments may have values added or removed. Add a +/- symbol in front of the argument. Add a secondary DNS server.

```
6. [root@desktopX ~]# nmcli con mod "static" +ipv4.dns
 8.8.8.8
```

7. Replace the static IP address and gateway.

```
8. [root@desktopX ~]# nmcli con mod "static"
 ipv4.addresses "172.25.X.10/24 172.25.X.254"
```

9. Add a secondary IP address without a gateway.

```
10. [root@desktopX ~]# nmcli con mod "static"
 +ipv4.addresses 10.10.10.10/16
```

**Important**



The **nmcli con mod** will save the setting to the configuration files. To activate the changes, the connection needs to be activated or reactivated.

```
[root@desktopX ~]# nmcli con up "static"
```

## Summary of **nmcli** commands

Basic device and connection commands for **nmcli**:

□

**nmcli** commands

| Command                                     | Use                                                                         |
|---------------------------------------------|-----------------------------------------------------------------------------|
| <code>nmcli dev status</code>               | List all devices.                                                           |
| <code>nmcli con show</code>                 | List all connections.                                                       |
| <code>nmcli con up "&lt;ID&gt;"</code>      | Activate a connection.                                                      |
| <code>nmcli con down "&lt;ID&gt;"</code>    | Deactivate a connection. The connection will restart if autoconnect is yes. |
| <code>nmcli dev dis &lt;DEV&gt;</code>      | Bring down an interface and temporarily disable autoconnect.                |
| <code>nmcli net off</code>                  | Disable all managed interfaces.                                             |
| <code>nmcli con add ...</code>              | Add a new connection.                                                       |
| <code>nmcli con mod "&lt;ID&gt;" ...</code> | Modify a connection.                                                        |
| <code>nmcli con del "&lt;ID&gt;"</code>     | Delete a connection.                                                        |

### Note

The **nmcli** command also has an interactive edit mode. For a graphical interface, use **nm-connection-editor**.

### References

**nmcli**(1), **nmcli-examples**(5), and **nm-settings**(5) man pages

Additional information may be available in the section on using the NetworkManager command line tool **nmcli** in the *Red Hat Enterprise Linux Networking Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

## 11.6. Practice: Configuring Networking with **nmcli**

### Guided exercise

In this lab, you will configure network settings using **nmcli**.

#### Outcomes:

Convert a system from DHCP to static configuration.

#### Before you begin...

Reset your serverX system.

1. View network settings using **nmcli**.

- a. Show all connections.

```
b. [student@serverX ~]$ nmcli con show
c. NAME UUID TYPE DEVICE
d. System eth0 5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03 802-3-ethernet
 eth0
```

e. Display all configuration settings for the active connection.

```
f. [student@serverX ~]$ nmcli con show "System eth0"
g. connection.id: System eth0
h. connection.uuid:
 5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03
i. connection.interface-name: eth0
j. connection.type: 802-3-ethernet
k. connection.autoconnect: yes
l. connection.timestamp: 1394813303
m. connection.read-only: no
n. connection.permissions:
o. ...
p. IP4.ADDRESS[1]: ip = 172.25.X.11/16, gw =
 172.25.0.1
q. IP4.DNS[1]: 172.25.254.254
r. IP4.DOMAIN[1]: example.com
s. ...
```

t. Show device status.

```
u. [student@serverX ~]$ nmcli dev status
v. DEVICE TYPE STATE CONNECTION
w. eth0 ethernet connected System eth0
x. lo loopback unmanaged --
```

y. Display the settings for the eth0 device.

```
z. [student@serverX ~]$ nmcli dev show eth0
aa. GENERAL.DEVICE: eth0
bb. GENERAL.TYPE: ethernet
cc. GENERAL.HWADDR: 52:54:00:00:00:0B
dd. GENERAL.MTU: 1500
ee. GENERAL.STATE: 100 (connected)
```

```

ff. GENERAL.CONNECTION: System eth0
gg. GENERAL.CON-PATH:
 /org/freedesktop/NetworkManager/ActiveConnection/1
hh. WIRED-PROPERTIES.CARRIER: on
ii. IP4.ADDRESS[1]: ip = 172.25.X.11/16, gw =
 172.25.0.1
jj. IP4.DNS[1]: 172.25.254.254
kk. IP4.DOMAIN[1]: example.com
ll. IP6.ADDRESS[1]: ip = fe80::5054:ff:fe00:b/64, gw
 = ::

```

2. Create a static connection with the same IPv4 address, network prefix, and default gateway. Name the new connection *static-eth0*.

#### Warning

Since access to your machine is provided over the primary network connection, setting incorrect values during network configuration may make your machine unreachable. If this happens, use the **Reset** button located above what used to be your machine's graphical display and try again.

```

[student@serverX ~]$ sudo nmcli con add con-name
"static-eth0" ifname eth0 type ethernet ip4
172.25.X.11/16 gw4 172.25.0.1

Connection 'static-eth0' (f3e8dd32-3c9d-48f6-9066-551e5b6e612d) successfully
added.

```

3. Modify the new connection to add the DNS setting.

```

4. [student@serverX ~]$ sudo nmcli con mod "static-eth0"
ip4.dns 172.25.254.254

```

5. Display and activate the new connection.

- a. View all connections.

```

b. [student@serverX ~]$ nmcli con show
c. NAME UUID TYPE DEVICE
d. static-eth0 f3e8dd32-3c9d-48f6-9066-551e5b6e612d 802-3-ethernet
 --
e. System eth0 5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03 802-3-ethernet
 eth0

```

- f. View the active connection.

```

g. [student@serverX ~]$ nmcli con show --active

```

```
h. System eth0 5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03 802-3-ethernet
 eth0
```

i. Activate the new connection.

```
j. [student@serverX ~]$ sudo nmcli con up
 "static-eth0"
```

```
k. Connection successfully activated (D-Bus active path:
 /org/freedesktop/NetworkManager/ActiveConnection/3)
```

l. View the active connection.

```
m. [student@serverX ~]$ nmcli con show --active
```

```
n. NAME UUID TYPE DEVICE
o. static-eth0 f3e8dd32-3c9d-48f6-9066-551e5b6e612d 802-3-ethernet
 eth0
```

6. Test the connectivity using the new network addresses.

a. Verify the IP address.

```
b. [student@serverX ~]$ ip addr show eth0
```

```
c. 2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast
 state UP qlen 1000
```

```
d. link/ether 52:54:00:00:00:0b brd ff:ff:ff:ff:ff:ff
```

```
e. inet 172.25.X.11/16 brd 172.25.255.255 scope global eth0
```

```
f. valid_lft forever preferred_lft forever
```

```
g. inet6 fe80::5054:ff:fe00:b/64 scope link
```

```
h. valid_lft forever preferred_lft forever
```

i. Verify the default gateway.

```
j. [student@serverX ~]$ ip route
```

```
k. default via 172.25.0.1 dev eth0 proto static metric 1024
```

```
l. 172.25.0.0/16 dev eth0 proto kernel scope link src 172.25.X.11
```

m. Ping the DNS address.

```
n. [student@serverX ~]$ ping -c3 172.25.254.254
```

```
o. PING 172.25.254.254 (172.25.254.254) 56(84) bytes of data.
```

```
p. 64 bytes from 172.25.254.254: icmp_seq=1 ttl=64 time=0.419 ms
```

```
q. 64 bytes from 172.25.254.254: icmp_seq=2 ttl=64 time=0.598 ms
```

```
r. 64 bytes from 172.25.254.254: icmp_seq=3 ttl=64 time=0.503 ms
```

```
s.
```

```

t. --- 172.25.254.254 ping statistics ---
u. 3 packets transmitted, 3 received, 0% packet loss, time 1999ms
v. rtt min/avg/max/mdev = 0.419/0.506/0.598/0.077 ms

```

7. Configure the original connection so that it does not start at boot and verify that the static connection is used when the system reboots.

- a. Disable the original connection from autostarting at boot.

```

b. [student@serverX ~]$ sudo nmcli con mod "System
 eth0" \
c. > connection.autoconnect no

```

- d. Reboot the system.

```

e. [student@serverX ~]$ reboot

```

- f. View the active connection.

```

g. [student@serverX ~]$ nmcli con show --active
h. NAME UUID TYPE DEVICE

```

```

 static-eth0 f3e8dd32-3c9d-48f6-9066-551e5b6e612d 802-3-ethernet eth0

```

## 11.7. Editing Network Configuration Files

### Objectives

After completing this section, students should be able to modify network settings by editing the configuration files.

### Modifying network configuration

It is also possible to configure the network by editing interface configuration files. Interface configuration files control the software interfaces for individual network devices. These files are usually named `/etc/sysconfig/network-scripts/ifcfg-<name>`, where `<name>` refers to the name of the device or connection that the configuration file controls. The following are standard variables found in the file used for static or dynamic configuration.

□

Configuration Options for `ifcfg` File

| <i>Static</i>                                                                | <i>Dynamic</i> | <i>Either</i>                                                         |
|------------------------------------------------------------------------------|----------------|-----------------------------------------------------------------------|
| BOOTPROTO=none<br>IPADDR0=172.25.X.10<br>PREFIX0=24<br>GATEWAY0=172.25.X.254 | BOOTPROTO=dhcp | DEVICE=eth0<br>NAME="System eth0"<br>ONBOOT=yes<br>UUID=f3e8dd32-3... |

| <i>Static</i>                                     | <i>Dynamic</i> | <i>Either</i>      |
|---------------------------------------------------|----------------|--------------------|
| <b>DEFROUTE=yes</b><br><b>DNS1=172.25.254.254</b> |                | <b>USERCTL=yes</b> |

In the static settings, variables for IP address, prefix, and gateway have a number at the end. **This allows multiple sets of values to be assigned to the interface.** The DNS variable also has a number which is used to specify the order of lookup when multiple servers are specified.

After modifying the configuration files, run **nmcli con reload** to make NetworkManager read the configuration changes. The interface still needs to be restarted for changes to take effect.

```
[root@serverX ~]# nmcli con reload
[root@serverX ~]# nmcli con down "System eth0"
[root@serverX ~]# nmcli con up "System eth0"
```

### References

**nmcli**(1) man page

Additional information may be available in the chapter on configuring networking in the *Red Hat Enterprise Linux Networking Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

## 11.8. Practice: Editing Network Configuration Files

### Guided exercise

In this lab, you will edit network configuration files.

#### Outcomes:

An additional network address added to each system.

#### Before you begin...

Reset your serverX and desktopX systems.

1. As the root user, edit the **/etc/sysconfig/network-scripts/ifcfg-eth0** on serverX to add an additional address of **10.0.X.1/24**.

- a. Append an entry to the file to specify the IPv4 address.

```
b. [root@serverX ~]# echo "IPADDR1=10.0.X.1" >>
/etc/sysconfig/network-scripts/ifcfg-eth0
```

- c. Append an entry to the file to specify the network prefix.

```
d. [root@serverX ~]# echo "PREFIX1=24" >>
/etc/sysconfig/network-scripts/ifcfg-eth0
```

2. Activate the new address.

- a. Reload the configuration changes.

```
b. [root@serverX ~]# nmcli con reload
```

- c. Restart the connection with the new settings.

```
d. [root@serverX ~]# nmcli con up "System eth0"
```

3. As the root user, edit the `/etc/sysconfig/network-scripts/ifcfg-eth0` on desktopX to add an additional address of `10.0.X.2/24` and load the new configuration.

- a. Modify the file to add the IPv4 and network prefix.

```
b. [root@desktopX ~]# echo "IPADDR1=10.0.X.2" >>
/etc/sysconfig/network-scripts/ifcfg-eth0
```

```
c. [root@desktopX ~]# echo "PREFIX1=24" >>
/etc/sysconfig/network-scripts/ifcfg-eth0
```

- d. Reload the configuration changes.

```
e. [root@desktopX ~]# nmcli con reload
```

- f. Bring up the connection with the new settings.

```
g. [root@desktopX ~]# nmcli con up "System eth0"
```

4. Test the connectivity using the new network addresses.

- a. On serverX, verify the IP address.

```
b. [root@serverX ~]# ip addr
```

- c. On serverX, ping the new address of desktopX.

```
d. [root@serverX ~]# ping 10.0.X.2
```

- e. On desktopX, verify the IP address.

```
f. [root@desktopX ~]# ip addr
```

- g. On desktopX, ping the new address of serverX.

```
h. [root@desktopX ~]# ping 10.0.X.1
```

## 11.9. Configuring Host Names and Name Resolution

### Objectives

After completing this section, students should be able to configure and test system host name and name resolution.

### Changing the system host name

*Configuring host names and name resolution*

The **hostname** command displays or temporarily modifies the system's fully qualified host name.

```
[root@desktopX ~]# hostname
desktopX.example.com
```

A static host name may be specified in the **/etc/hostname** file.

The **hostnamectl** command is used to **modify this file** and may be used to view the status of the system's fully qualified host name. If this file does not exist, the host name is set by a reverse DNS query once the interface has an IP address assigned.

```
[root@desktopX ~]# hostnamectl set-hostname
desktopX.example.com
[root@desktopX ~]# hostnamectl status
Static hostname: desktopX.example.com
 Icon name: computer
 Chassis: n/a
 Machine ID: 9f6fb63045a845d79e5e870b914c61c9
 Boot ID: aa6c3259825e4b8c92bd0f601089ddf7
 Virtualization: kvm
 Operating System: Red Hat Enterprise Linux Server 7.0 (Maipo)
 CPE OS Name: cpe:/o:redhat:enterprise_linux:7.0:beta:server
 Kernel: Linux 3.10.0-97.el7.x86_64
 Architecture: x86_64
[root@desktopX ~]# cat /etc/hostname
desktopX.example.com
```

### Important

The static host name is stored in **/etc/hostname**. Previous versions of Red Hat Enterprise Linux stored the host name as a variable in the **/etc/sysconfig/network** file.

## Configuring name resolution

The *stub resolver* is used to convert host names to IP addresses or the reverse. The contents of the file **/etc/hosts** are checked first.

```
[root@desktopX ~]# cat /etc/hosts
127.0.0.1 localhost localhost.localdomain localhost4
localhost4.localhostdomain4
::1 localhost localhost.localdomain localhost6
localhost6.localhostdomain6

172.25.254.254 classroom.example.com
172.25.254.254 content.example.com
```



The `getent hosts hostname` command can be used to test host name resolution with the `/etc/hosts` file.

If an entry is not found in that file, the stub resolver looks for the information from a DNS nameserver.

The `/etc/resolv.conf` file controls how this query is done:

- **nameserver**: the IP address of a nameserver to query. Up to three `nameserver` directives may be given to provide backups if one is down.
- **search**: a list of domain names to try with a short host name. Both this and **domain** should not be set in the same file; if they are, the last instance wins.

See `resolv.conf(5)` for details.

```
[root@desktopX ~]# cat /etc/resolv.conf
Generated by NetworkManager
domain example.com
search example.com
nameserver 172.25.254.254
```

NetworkManager will update the `/etc/resolv.conf` file using DNS settings in the connection configuration files. Use the `nmcli` to modify the connections.

```
[root@desktopX ~]# nmcli con mod ID ipv4.dns IP
[root@desktopX ~]# nmcli con down ID
[root@desktopX ~]# nmcli con up ID
[root@desktopX ~]# cat
/etc/sysconfig/network-scripts/ifcfg-ID
...
DNS1=8.8.8.8
...
```

The default behavior of `nmcli con mod ID ipv4.dns IP` is to replace any previous DNS settings with the new IP list provided. A `+/-` symbol in front of the `ipv4.dns` argument will add or remove an individual entry.

```
[root@desktopX ~]# nmcli con mod ID +ipv4.dns IP
```

The `host HOSTNAME` command can be used to test DNS server connectivity.

```
[root@desktopX ~]# host classroom.example.com
classroom.example.com has address 172.25.254.254

[root@desktopX ~]# host 172.25.254.254
254.254.25.172.in-addr.arpa domain name pointer classroom.example.com.
```

Important

If DHCP is in use, `/etc/resolv.conf` is automatically rewritten as interfaces are started, unless you specify `PEERDNS=no` in the relevant interface configuration files. The change can be made with `nmcli`.

```
[root@desktopX ~]# nmcli con mod "System eth0"
ipv4.ignore-auto-dns yes
```

## References

`nmcli(1)`, `hostnamectl(1)`, `hosts(5)`, `getent(1)`, `host(1)`, and `resolv.conf(5)` man pages

Additional information may be available in the chapter on configuring host names in the *Red Hat Enterprise Linux Networking Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

## 11.10. Practice: Configuring Host Names and Name

### Resolution

#### Guided exercise

In this lab, you will configure the system host name and name resolution.

#### Outcomes:

Customized host name and name resolution settings.

#### Before you begin...

Reset your serverX system.

1. View the current host name settings.

- a. Display the current host name.

```
b. [student@serverX ~]$ hostname
c. serverX.example.com
```

- d. Display the host name status.

```
e. [student@serverX ~]$ hostnamectl status
f. Static hostname: n/a
g. Transient hostname: serverX.example.com
h. Icon name: computer
i. Chassis: n/a
j. Machine ID: 9f6fb63045a845d79e5e870b914c61c9
k. Boot ID: d4ec3a2e8d3c48749aa82738c0ea946a
l. Operating System: Red Hat Enterprise Linux Server 7.0 (Maipo)
m. CPE OS Name: cpe:/o:redhat:enterprise_linux:7.0:beta:server
n. Kernel: Linux 3.10.0-97.el7.x86_64
o. Architecture: x86_64
```

2. Set a static host name to match the current transient host name.

- a. Change the host name and host name configuration file. Replace the *x* with your station number and match the output of the previous step.

```
b. [student@serverX ~]$ sudo hostnamectl
set-hostname serverX.example.com
```

- c. View the configuration file providing the host name at network start.

```
d. [student@serverX ~]$ cat /etc/hostname
e. serverX.example.com
```

- f. Display the host name status.

```
g. [student@serverX ~]$ hostnamectl status
h. Static hostname: serverX.example.com
i. Icon name: computer
j. Chassis: n/a
k. Machine ID: 9f6fb63045a845d79e5e870b914c61c9
l. Boot ID: d4ec3a2e8d3c48749aa82738c0ea946a
m. Operating System: Red Hat Enterprise Linux Server 7.0 (Maipo)
n. CPE OS Name: cpe:/o:redhat:enterprise_linux:7.0:beta:server
o. Kernel: Linux 3.10.0-97.el7.x86_64
p. Architecture: x86_64
```

3. Temporarily change the host name.

- a. Change the host name.

```
b. [student@serverX ~]$ sudo hostname testname
```

- c. Display the current host name.

```
d. [student@serverX ~]$ hostname
e. testname
```

- f. View the configuration file providing the host name at network start.

```
g. [student@serverX ~]$ cat /etc/hostname
h. serverX.example.com
```

- i. Reboot the system.

```
j. [student@serverX ~]$ reboot
```

- k. Display the current host name.

```
l. [student@serverX ~]$ hostname
m. serverX.example.com
```

4. Add a local **nickname** for the classroom server.

绰号

- a. Look up the IP address of the classroom.example.com.

```
b. [student@serverX ~]$ host
 classroom.example.com
c. classroom.example.com has address 172.25.254.254
```

- d. Modify **/etc/hosts** so that the name **class** has the IP address 172.25.254.254 and can be used to communicate with classroom.example.com.

```
e. [student@serverX ~]$ sudo vi /etc/hosts
f. [student@serverX ~]$ cat /etc/hosts
g. 127.0.0.1 localhost localhost.localdomain localhost4
 localhost4.localdomain4
h. ::1 localhost localhost.localdomain localhost6
 localhost6.localdomain6
i.
j. 172.25.254.254 classroom.example.com class
k. 172.25.254.254 content.example.com
```

- l. Look up the IP address of the class.

```
m. [student@serverX ~]$ host class
n. Host class not found: 2(SERVFAIL)
o. [student@serverX ~]$ getent hosts class
p. 172.25.254.254 classroom.example.com class
```

已接受

- q. Ping class.

```
r. [student@serverX ~]$ ping -c3 class
s. PING classroom.example.com (172.25.254.254) 56(84) bytes of data.
t. 64 bytes from classroom.example.com (172.25.254.254): icmp_seq=1
 ttl=64 time=0.397 ms
u. 64 bytes from classroom.example.com (172.25.254.254): icmp_seq=2
 ttl=64 time=0.447 ms
v. 64 bytes from classroom.example.com (172.25.254.254): icmp_seq=3
 ttl=64 time=0.470 ms
w.
x. --- classroom.example.com ping statistics ---
```

```
y. 3 packets transmitted, 3 received, 0% packet loss, time 2000ms
```

```
rtt min/avg/max/mdev = 0.397/0.438/0.470/0.030 ms
```

## 11.11. Lab: Managing Red Hat Enterprise Linux Networking

### Performance checklist

In this lab, you will configure basic IPv4 networking on Red Hat Enterprise Linux systems.

#### Outcomes:

The primary interface has two static IPv4 addresses configured.

#### Before you begin...

Reset your desktopX system.

1. Create a new connection with a static network connection using the settings in the table. Be sure to replace the *x* with the correct number for your systems.

| Parameter       | Setting        |
|-----------------|----------------|
| Connection name | lab            |
| IP address      | 172.25.x.10/16 |
| Gateway address | 172.25.0.1     |
| DNS address     | 172.25.254.254 |

2. 

```
[root@desktopX ~]# nmcli con add con-name lab ifname eth0 type ethernet ip4 172.25.x.10/16 gw4 172.25.0.1
```
3. 

```
[root@desktopX ~]# nmcli con mod "lab" ipv4.dns 172.25.254.254
```

4. Configure the new connection to be autostarted. Other connections should not start automatically.

5. 

```
[root@desktopX ~]# nmcli con mod "lab" connection.autoconnect yes
```
6. 

```
[root@desktopX ~]# nmcli con mod "System eth0" connection.autoconnect no
```

7. Modify the new connection so that it also uses the address 10.0.X.1/24.

8. 

```
[root@desktopX ~]# nmcli con mod "lab" +ipv4.addresses 10.0.x.1/24
```

Or alternately:

```
[root@desktopX ~]# echo "IPADDR1=10.0.X.1" >>
/etc/sysconfig/network-scripts/ifcfg-lab

[root@desktopX ~]# echo "PREFIX1=24" >>
/etc/sysconfig/network-scripts/ifcfg-lab
```

9. Configure the **hosts** file so that 10.0.X.1 can be referenced as "private".

```
10. [root@desktopX ~]# echo "10.0.X.1 private" >>
/etc/hosts
```

11. Reboot the system, then run **lab network grade** to verify settings.

```
12. [root@desktopX ~]# lab network grade
```

# Chapter 12. Archiving and Copying Files Between Systems

12.1. Managing Compressed tar Archives

12.2. Practice: Backing Up and Restoring Files From a tar Archive

12.3. Copying Files Between Systems Securely

12.4. Practice: Copying Files Over the Network With `scp`

12.5. Synchronizing Files Between Systems Securely

12.6. Practice: Synchronizing Two Directories Securely With `rsync`

12.7. Lab: Archiving and Copying Files Between Systems

**Introduction**

## Overview:

|                   |                                                                                                                                                                                                                                                  |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Goal</b>       | To archive and copy files from one system to another.                                                                                                                                                                                            |
| <b>Objectives</b> | <ul style="list-style-type: none"><li>○ Use <code>tar</code> to create new compressed archive files and extract files from existing archive files.</li><li>○ Copy files securely to or from a remote system running <code>sshd</code>.</li></ul> |

| Overview: |                                                                                                                                                                                                                                       |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|           | <ul style="list-style-type: none"> <li>Securely synchronize the contents of a local file or directory with a remote copy.</li> </ul>                                                                                                  |
| Sections  | <ul style="list-style-type: none"> <li>Managing Compressed tar Archives (and Practice)</li> <li>Copying Files Between Systems Securely (and Practice)</li> <li>Synchronizing Files Between Systems Securely (and Practice)</li> </ul> |
| Lab       | <ul style="list-style-type: none"> <li>Archiving and Copying Files Between Systems</li> </ul>                                                                                                                                         |

## 12.1. Managing Compressed tar Archives

### Objective

After completing this section, students should be able to use tar to create new compressed archive files and extract files from existing archive files.

*Archives and compression*

### What is tar?

Archiving and compressing files are useful when creating backups and transferring data across a network. One of the oldest and most common commands for creating and working with backup archives is the **tar** command.

With **tar**, users can gather large sets of files into a single file (archive). The archive can be compressed using **gzip**, **bzip2**, or **xz** compression.

The tar command can list the contents of archives or extract their files to the current system. Examples of how to use the **tar** command are included in this section.

### Operate the tar command

To use the **tar** command, one of the three following actions is required:

- c** (create an archive)
- t** (list the contents of an archive)
- x** (extract an archive)

Commonly used options are:

- f file name** (file name of the archive to operate on)
- v** (verbosity; useful to see which files get added to or extracted from the archive)

#### Note

A leading **-** is not required for tar options.

### Archive files and directories with tar

Before creating a tar archive, verify that there is no other archive in the directory with the same name as the new archive to be created. The **tar** command will overwrite an existing archive without any feedback.

The first option to use when creating a new archive is the **c** option, followed by the **f** option, then a single space, then the file name of the archive to be created, and finally the list of files and directories that should get added to the archive. The archive is created in the current directory unless specified otherwise.

In the following example, an archive named **archive.tar** is created with the contents of file1, file2, and file3 in the user's home directory.

```
[user@host ~]# tar cf archive.tar file1 file2 file3
[user@host ~]# ls archive.tar
archive.tar
```

#### Note

When archiving files by absolute path names, the leading `/` of the path is removed from the file name by default. This helps avoid mistakes which could cause important files to be overwritten. Files are normally extracted relative to the current working directory of the `tar` command.

For `tar` to be able to archive the selected files, it is mandatory that the user executing the `tar` command is able to read the file(s). For example, creating a new archive of the `/etc` folder and all of its content requires root privileges, because only root is allowed to read all of the files there. An unprivileged user could create an archive of the `/etc` folder, but the archive would omit files which do not include read permission for the user and it would omit directories which do not include both read and execute permission for the user.

Create the tar archive `/root/etc.tar` with the `/etc` directory as content as user root:

```
[root@host ~]# tar cf /root/etc.tar /etc
tar: Removing leading `/' from member names
[root@host ~]#
```

#### Important

While `tar` stores ownership and permissions of the files, there are other attributes that are not stored in the `tar` archive by default, such as the SELinux context and ACLs. To store those extended attributes in the `tar` archive, the `--xattrs` option is required when creating an archive.

#### List contents of a tar archive

To list the content of an archive, the `t` and `f` options, accompanied by the archive to operate, are required.

List the content of the archive `/root/etc.tar`:

```
[root@host ~]# tar tf /root/etc.tar
etc/
etc/fstab
etc/crypttab
etc/mtab
...
```

#### Extract an archive created with tar

A tar archive should normally be extracted in an empty directory to ensure it does not overwrite any existing files. If files are extracted by root, `tar` attempts to preserve the original user and group ownership of the files. If a regular user extracts files using `tar`, the extracted files are owned by that user.

Extract the archive `/root/etc.tar` to the `/root/etcbakup` directory:

```
[root@host ~]# mkdir /root/etcbakup
```



```
[root@host ~]# cd /root/etcbakup
[root@host etcbakup]# tar xf /root/etc.tar
```

By default, when files get extracted from an archive, the umask is subtracted from the permissions of archive content. This is a security measure and prevents extracted regular files from receiving execute permissions by default. To preserve the permissions of an archived file, the **p** option is to be used when extracting an archive.

Extract the archive `/root/myscripts.tar` to the `/root/scripts` directory while preserving the permissions of the extracted files:

```
[root@host ~]# mkdir /root/scripts
[root@host ~]# cd /root/scripts
[root@host scripts]# tar xpf /root/myscripts.tar
```

### ❏ Create a compressed tar archive

There are three different compression methods supported by the **tar** command. The **gzip** compression is the fastest and oldest one, and is most widely available. The **bzip2** compression usually leads to smaller archive files compared to **gzip** and is less widely available than **gzip**, while the **xz** compression method is relatively new, but usually offers the best compression ratio of the methods available.

#### Note

The effectiveness of any compression algorithm depends on the exact nature of the data being compressed. Data files that are already compressed, such as compressed picture formats or rpm files, usually lead to a low compression ratio.

It is good practice to use a single top-level directory, which can contain other directories and files, to simplify extraction of the files in an organized way.

To create a compressed tar archive, one of the following **tar** options can be specified:

- **z** for gzip compression (filename.tar.gz or filename.tgz)
- **j** for bzip2 compression (filename.tar.bz2)
- **J** for xz compression (filename.tar.xz)

Create (c option) a gzip-compressed (z option) tar archive `/root/etcbakup.tar.gz` of the `/etc` directory on serverX:

```
[root@serverX ~]$ tar czf /root/etcbakup.tar.gz /etc
```

Create (c option) a bzip2-compressed (j option) tar archive `/root/logbackup.tar.bz2` of the `/var/log` directory on serverX:

```
[root@serverX ~]$ tar cjf /root/logbackup.tar.bz2 /var/log
```

Create (c option) a xz-compressed (J option) tar archive `/root/sshconfig.tar.xz` of the `/etc/ssh` directory on serverX:

```
[root@serverX ~]$ tar cJf /root/sshconfig.tar.xz /etc/ssh
```

## Extract a compressed tar archive

The first step when extracting a compressed tar archive is to determine where the archived files should be extracted to, then create and change to the target directory. To successfully extract the archive, it is usually not necessary to use the same compression option used when creating the archive, as the **tar** command will determine which compression was used. It is valid to add the decompression method to the **tar** options as follows:

Extract (x option) the contents of a gzip-compressed (z option) tar archive named **/root/etcbackup.tar.gz** to the directory **/tmp/etcbackup**:

```
[root@serverX ~]$ mkdir /tmp/etcbackup
[root@serverX ~]$ cd /tmp/etcbackup
[root@serverX etcbackup]$ tar xzf /root/etcbackup.tar.gz
```

Extract (x option) the contents of a bzip2-compressed (j option) tar archive named **/root/logbackup.tar.bz2** to the directory **/tmp/logbackup**:

```
[root@serverX ~]$ mkdir /tmp/logbackup
[root@serverX ~]$ cd /tmp/logbackup
[root@serverX logbackup]# tar xjf /root/logbackup.tar.bz2
```

Extract (x option) the contents of a xz-compressed (J option) tar archive named **/root/sshbackup.tar.xz** to the directory **/tmp/sshbackup**:

```
[root@serverX ~]$ mkdir /tmp/sshbackup
[root@serverX ~]$ cd /tmp/sshbackup
[root@serverX sshbackup]# tar xJf /root/sshbackup.tar.xz
```

### Note

Listing a compressed **tar** archive works in the same way as listing an uncompressed **tar** archive.

### Note

Additionally, **gzip**, **bzip2**, and **xz** can be used independently to compress single files. For example, **gzip etc.tar** results in the compressed file **etc.tar.gz**, while **bzip2 abc.tar** results in the compressed file **abc.tar.bz2** and **xz myarchive.tar** results in the compressed file **myarchive.tar.xz**.

The corresponding decompress commands are **gunzip**, **bunzip2**, and **unxz**. For example, **gunzip /tmp/etc.tar.gz** results in the uncompressed tar file **etc.tar**, while **bunzip2 abc.tar.bz2** results in the uncompressed tar file **abc.tar** and **unxz myarchive.tar.xz** results in the uncompressed tar file **myarchive.tar**.

## Overview of tar options

The **tar** command has many options to use. The following table lists some common options and their meanings.

□

Overview of **tar** options

| Option | Meaning |
|--------|---------|
|--------|---------|

| Option   | Meaning                                                                                                                         |
|----------|---------------------------------------------------------------------------------------------------------------------------------|
| <b>c</b> | Create a new archive.                                                                                                           |
| <b>x</b> | Extract from an existing archive.                                                                                               |
| <b>t</b> | List the contents of an archive.                                                                                                |
| <b>v</b> | Verbose; shows which files get archived or extracted.                                                                           |
| <b>f</b> | File name; this option needs to be followed by the file name of the archive to use/create.                                      |
| <b>p</b> | Preserve the permissions of files and directories when extracting an archive, without subtracting the umask.                    |
| <b>z</b> | Use <b>gzip</b> compression ( <b>.tar.gz</b> ).                                                                                 |
| <b>j</b> | Use <b>bzip2</b> compression ( <b>.tar.bz2</b> ). <b>bzip2</b> typically achieves a better compression ratio than <b>gzip</b> . |
| <b>J</b> | Use <b>xz</b> compression ( <b>.tar.xz</b> ). <b>xz</b> typically achieves a better compression ratio than <b>bzip2</b> .       |

#### References

[tar\(1\)](#), [gzip\(1\)](#), [gunzip\(1\)](#), [bzip2\(1\)](#), [bunzip2\(1\)](#), [xz\(1\)](#), [unxz\(1\)](#) man pages

## 12.2. Practice: Backing Up and Restoring Files From a tar

### Archive

#### Guided exercise

In this lab, students will create and extract archives with **tar**.

#### Outcomes:

Students will back up a directory tree and extract the archive content to another location.

1. Since only the root user can read all the contents of the **/etc** directory, we are going to back up the directory and log into serverX as **root**.

```
[student@desktopX ~]$ ssh root@serverX
```

2. Create an archive of **/etc** using **gzip** compression to back up the configuration file directory **/etc**. Save the file as **/tmp/etc.tar.gz**.

```
[root@serverX ~]# tar czf /tmp/etc.tar.gz /etc
```

3. Verify that the backup file **etc.tar.gz** is a valid archive by decompressing the file to a newly created directory named **/backuptest** on serverX.
  - a. Create the target directory **/backuptest**.

```
[root@serverX ~]# mkdir /backuptest
```

- b. Switch to the `/backuptest` directory, where we will extract the files from the `etc.tar.gz` archive.

```
[root@serverX ~]# cd /backuptest
```

- c. Extract the `etc.tar.gz` archive to the `/backuptest` directory.

```
[root@serverX backuptest]# tar xzf
/tmp/etc.tar.gz
```

## 12.3. Copying Files Between Systems Securely

### Objectives

After completing this section, students should be able to copy files securely to or from a remote system running `sshd`.

### Copy files to or from a remote location with `scp`

*Copying files between systems securely*

远程连接Linux最常用的  
安全方式-SSH

The `ssh` command is useful for securely running shell commands on remote systems. It can also be used to securely copy files from one machine to another. The `scp` command transfers files from a remote host to the local system or from the local system to a remote host. It utilizes the SSH server for authentication and encrypted data transfer.

SSH认证和加密数据传输

Remote file system locations are always specified in the format `[user@]host:/path` for either the source or target location of the files to be transferred. The `user@` portion is optional and, if it is missing, the current local user that invokes the `scp` command is used. Before the transfer is initiated, the user must authenticate with the SSH server by password or SSH keys.

部分

This example shows how to copy the local files on desktopX, `/etc/yum.conf` and `/etc/hosts`, securely to the account student on the remote system serverX into the directory `/home/student/`:

```
[student@desktopX ~]$ scp /etc/yum.conf /etc/hosts
serverX:/home/student

student@serverX's password: student
yum.conf 100% 813 0.8KB/s 00:00
hosts 100% 227 0.2KB/s 00:00
```

A user can copy a file from a remote account on a remote machine to the local file system with `scp`. In this example, copy the file `/etc/hostname` from the account student on the serverX machine to the local directory `/home/student/`.

```
[student@desktopX ~]$ scp serverX:/etc/hostname
/home/student/
student@serverX's password: student
```

安全地拷贝文件

```
hostname 100% 22 0.0KB/s 00:00
```

To copy a whole directory tree recursively, the `-r` option is available. In the following example, the remote directory `/var/log` on serverX is copied recursively to the local directory `/tmp/` on desktopX. To be able to read all the contents of the `/etc` directory, the root user must connect to the remote location.

```
[student@desktopX ~]$ scp -r root@serverX:/var/log /tmp
root@serverX's password: redhat
...
```

## Transfer files remotely with `sftp`

更好的

If an interactive tool is preferred when uploading or downloading files to a SSH server, the `sftp` command can be used. A session with `sftp` is similar to a classic FTP session, but uses the secure authentication mechanism and encrypted data transfer of the SSH server.

To initiate a `sftp` session, the `sftp` expects a remote location in the format `[user@]host`, where the `user@` portion is optional and, if it is missing, the user invoking the `sftp` command is used. To establish the `sftp` session, authenticating with any of the methods the SSH server accepts is necessary.

```
[student@desktopX ~]$ sftp serverX
student@serverX's password: student
Connected to serverX.
sftp>
```

The `sftp` session accepts various commands that work the same way on the remote file system as they do in the local file system, such as `ls`, `cd`, `mkdir`, `rmdir`, and `pwd`. In addition, there are the `put` and `get` commands for uploading and downloading files. The `exit` command exits the `sftp` session.

Upload the local file `/etc/hosts` to the newly created directory `/home/student/hostbackup` on the remote host serverX. The `sftp` session always assumes that the `put` command is followed by a file on the local file system and starts in the connecting user's home directory; in this case, `/home/student`:

```
sftp> mkdir hostbackup
sftp> cd hostbackup
sftp> put /etc/hosts
Uploading /etc/hosts to /home/student/hostbackup/hosts
/etc/hosts 100% 227 0.2KB/s 00:00
sftp>
```

To download the remote file `/etc/yum.conf` from the remote host to the current directory on the local file system, execute the command `get /etc/yum.conf` and exit the `sftp` session with the `exit` command.

```
sftp> get /etc/yum.conf
Fetching /etc/yum.conf to yum.conf
/etc/yum.conf 100% 813 0.8KB/s 00:00
sftp> exit
[student@desktopX ~]$
```

## References

`scp(1)` and `sftp(1)` man pages

## 12.4. Practice: Copying Files Over the Network With `scp`

### Guided exercise

In this lab, students will copy files from a remote system to a local directory by using `scp`.

#### Outcomes:

Students will copy files from a remote host to a directory on the local machine.

- Remotely copy the `/etc/ssh` directory on the serverX machine to the newly created directory `/home/student/serverbackup` on desktopX by using `scp`.
  1. Create the target directory `/home/student/serverbackup` on desktopX.

```
[student@desktopX ~]$ mkdir
/home/student/serverbackup
```

2. Recursively copy the directory `/etc/ssh` from serverX to the `/home/student/serverbackup` directory on desktopX with the `scp` command. Note that only the root user can read all the content in the `/etc/ssh` directory.

```
[student@desktopX ~]$ scp -r
root@serverX:/etc/ssh /home/student/serverbackup
```

## 12.5. Synchronizing Files Between Systems Securely

### Objectives

After completing this section, students should be able to efficiently and securely synchronize the contents of a local file or directory with a remote copy.

### Synchronize files and folders with `rsync`

Remote File Transfers

[同步文件和文件夹](#)

The `rsync` tool is another way to securely copy files from one system to another. It differs from `scp` in that if two files or directories are similar between two systems, `rsync` only needs to copy the differences between the systems, while `scp` would need to copy everything.

One of the advantages of `rsync` is that it can copy files between a local system and a remote system securely and efficiently. While the initial synchronization of a directory takes about the same time as

copying it, any subsequent synchronization only requires the differences to be copied over the network.

并发的同步

测试

One of the most important options of **rsync** is the **-n option** to perform a dry run. A **dry run** is a

模拟

**simulation** of what happens when the command really gets executed. It will display the changes it will perform when the command is executed without the dry run option. It is recommended to perform a dry run of any **rsync** operation to ensure no important files get overwritten or deleted.

The two most common options when synchronizing files and folders with **rsync** are the **-a** and **-v** options. While the **-v** option adds verbosity to the output as the synchronization proceeds, the **-a** option stands for "archive mode" and enables the following options all in one:

- **-r**, synchronize recursively the whole directory tree
- **-l**, synchronize symbolic links
- **-p**, preserve permissions
- **-t**, preserve time stamps
- **-g**, preserve group ownership
- **-o**, preserve the owner of the files
- **-D**, synchronize device files

同步递归整个目录

详细信息

存档模式，递归并保留对象属性，等同于 -rlptgoD

保留硬链接文件

While the **-a** already synchronizes symbolic links, there are additional options necessary to preserve hardlinks, as they are treated as separate files instead. The **-H** option enables the handling of hardlinks, so the **rsync** command will identify the hardlinks present in the source folder and link the files accordingly in the destination folder instead of just copying them as separate files.

#### Note

The **-a** option does not synchronize advanced file permissions, such as ACLs or SELinux file contexts. To enable the synchronization of **ACLs**, the **-A** option is required in addition to the **-a** option, while to synchronize the **SELinux** contexts from the source files to the target files, the **-X** option is to be added.

The basic way of using **rsync** is to synchronize two local folders. In the following example, the **/var/log** directory gets a synchronized copy in the **/tmp** folder. The **log** directory with its content is created in the **/tmp** directory.

```
[student@desktopX ~]$ su -
Password: redhat
[root@desktopX ~]# rsync -av /var/log /tmp/
...
```

To only synchronize the content of a folder without newly creating the folder in the target directory, a trailing slash needs to be added at the end of the source directory. In this example, the **log** directory is not created in the **/tmp** folder. Only the content of the **/var/log/** directory is synchronized into the **/tmp** folder.

```
[root@desktopX ~]# rsync -av /var/log/ /tmp
...
```

Important

When entering the source directory for `rsync`, it is very important to remember that whether a trailing slash is present on the directory name matters. It will determine whether the *directory* or just the *contents of the directory* are synchronized to the target. Note: Tab-completion will automatically add a trailing slash to the end of directory names.

Similar to `scp`, the `rsync` command expects remote file system locations to be specified in the format `[user@]host:/path`. In case the optional `user@` portion is missing, the user invoking the `rsync` command is used to connect to the remote location. It is possible to use a remote location as either source or target. In the following example, the local folder `/var/log` gets a synchronized copy in the `/tmp` directory on the serverX machine. For `rsync` to synchronize ownership of the transferred files, the target location must be written as user `root`, so connect to the remote system serverX as the root user. The connecting user `root` must then authenticate with the SSH server by any of the accepted methods; for example, password or SSH keys.

```
[root@desktopX ~]$ rsync -av /var/log serverX:/tmp
root@serverX's password: redhat
...
```

In the same way, the remote folder `/var/log` on serverX can be synchronized to the local directory `/tmp` on desktopX:

```
[root@desktopX ~]$ rsync -av serverX:/var/log /tmp
root@serverX's password: redhat
...
```

## References

[rsync\(1\) man page](#)

## 12.6. Practice: Synchronizing Two Directories Securely

### With `rsync`

#### Guided exercise

In this lab, students will synchronize a folder with a remote system by using `rsync`.

#### Outcomes:

A directory will be synchronized from a remote machine to the local machine. After files on the remote machine have changed, it will be synchronized with the local machine again and only the modifications will be transferred. Both systems will end with identical content in the directories synchronized with `rsync`.

1. Securely create an initial copy of the `/var/log` directory tree on serverX to a newly created directory named `/serverlogs` on desktopX with the `rsync` command.
  - a. In order to create the target directory `/serverlogs`, switch to the root user account with the `su` command.

```
b. [student@desktopX ~]$ su -
c. Password: redhat
```



```
[root@desktopX ~]#
```

- d. Create the target directory **/serverlogs** on desktopX where the log files of serverX will be synchronized.

```
[root@desktopX ~]# mkdir /serverlogs
```

- e. Use the **rsync** command to synchronize the **/var/log** directory tree on serverX to the **/serverlogs** directory on desktopX. Note that only the root user can read all the content in the **/var/log** directory on serverX. All files will be transferred in the initial synchronization.

```
f. [root@desktopX ~]# rsync -av
root@serverX:/var/log /serverlogs
```

```
...
```

2. As root on serverX, run **logger "Log files synchronized"** to get a new entry in the log file **/var/log/messages** to reflect when the last synchronization took place.

```
3. [root@desktopX ~]# ssh root@serverX 'logger "Log files
synchronized" '
```

```
4. Password: redhat
```

```
[root@desktopX ~]#
```

5. Securely synchronize the **/var/log** directory tree on serverX to the **/serverlogs** directory on desktopX with the **rsync** command again. Note that this time only the changed log file(s) will be transferred.

```
6. [root@desktopX ~]# rsync -av root@serverX:/var/log
/serverlogs
```

```
...
```

## 12.7. Lab: Archiving and Copying Files Between Systems

### Performance checklist

In this lab, students will use **rsync**, **scp**, and **tar** to archive and back up folder contents.

#### Outcomes:

Students will synchronize a remote folder to a local directory; an archive is then created with the synchronized folder as content; the archive gets copied to the remote machine and extracted to a newly created directory.

#### Before you begin...

Reset your serverX system.

1. Synchronize the `/etc` directory tree on serverX to the `/conffigsync` directory on desktopX.
  - a. To be able to create the target directory `/conffigsync`, switch to the root user account using `su`.

```
b. [student@desktopX ~]$ su -
c. Password: redhat
```

```
[root@desktopX ~]#
```

- d. Create the target directory for the configuration files on desktopX.

```
[root@desktopX ~]# mkdir /conffigsync
```

- e. Use the `rsync` command to synchronize the `/etc` directory tree on serverX to the `/conffigsync` directory on desktopX. Be aware that only the root user can read all the content in the `/etc` directory on serverX.

```
f. [root@desktopX ~]# rsync -av root@serverX:/etc
/conffigsync
```

```
...
```

2. On desktopX, create an archive named `/root/configfile-backup-serverX.tar.gz` with the `/conffigsync` directory as content, and copy the archive to the `/root` directory on serverX for backup purposes with the `scp` command.

- a. Store the `/conffigsync` directory in the `/root/configfile-backup-serverX.tar.gz` archive.

```
[root@desktopX ~]# tar czf
/root/configfile-backup-serverX.tar.gz
/conffigsync
```

- b. Create a backup copy of the `/root/configfile-backup-serverX.tar.gz` on serverX

```
c. [root@desktopX conffigsync]# scp
/root/configfile-backup-serverX.tar.gz
root@serverX:/root
```

```
d. Password: redhat
```

```
...
```

3. To prepare the archived directory tree for comparison with the currently actively used configuration files on serverX, extract the contents of the `/root/configfile-backup-serverX.tar.gz` archive to the `/tmp/savedconfig/` directory on serverX.

- a. Connect to the serverX machine as root by using `ssh`.

```
b. [root@desktopX configsync]# ssh root@serverX
c. Password: redhat
```

```
[root@serverX ~]#
```

- d. Create the target directory `/tmp/savedconfig/`, where the contents of the `/root/configfile-backup-serverX.tar.gz` archive will be extracted.

```
[root@serverX ~]# mkdir /tmp/savedconfig
```

- e. Change to the target directory `/tmp/savedconfig/` on serverX.

```
f. [root@serverX ~]# cd /tmp/savedconfig
```

```
[root@serverX savedconfig]#
```

- g. Extract the contents of the `/root/configfile-backup-serverX.tar.gz` archive to the `/tmp/savedconfig/` directory on serverX.

```
[root@serverX savedconfig]# tar xzf
/root/configfile-backup-serverX.tar.gz
```

# Chapter 13. Installing and Updating Software Packages

13.1. Attaching Systems to Subscriptions for Software Updates

13.2. Practice: Red Hat Subscription Management

13.3. RPM Software Packages and Yum

13.4. Practice: RPM Software Packages

13.5. Managing Software Updates with yum

13.6. Practice: Installing and Updating Software with `yum`

13.7. Enabling `yum` Software Repositories

13.8. Practice: Enabling Software Repositories

13.9. Examining RPM Package Files

13.10. Practice: Working with RPM Package Files

13.11. Lab: Installing and Updating Software Packages

### Introduction

| Overview   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Goal       | To download, install, update, and manage software packages from Red Hat and YUM package repositories.                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Objectives | <ul style="list-style-type: none"><li>○ Register systems with your Red Hat account and entitle them to software updates for installed products.</li><li>○ Explain what an RPM package is and how RPM packages are used to manage software on a Red Hat Enterprise Linux system.</li><li>○ Find, install, and update software packages using the <code>yum</code> command.</li><li>○ Enable and disable use of Red Hat or third-party YUM repositories.</li><li>○ Examine and install downloaded software package files.</li></ul> |
| Sections   | <ul style="list-style-type: none"><li>○ Attaching Systems to Subscriptions for Software Updates (and Practice)</li><li>○ RPM Software Packages and YUM (and Practice)</li><li>○ Managing Software Updates with <code>yum</code> (and Practice)</li><li>○ Enabling <code>yum</code> Software Repositories (and Practice)</li><li>○ Examining RPM Package Files (and Practice)</li></ul>                                                                                                                                            |
| Lab        | <ul style="list-style-type: none"><li>○ Installing and Updating Software Packages</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## ▮ 13.1. Attaching Systems to Subscriptions for Software

### Updates

#### ▮ Objectives

Register systems with your Red Hat account and entitle them to software updates for installed products.

## Red Hat Subscription Management

Red Hat Subscription Management provides tools that can be used to entitle machines to product subscriptions, allowing administrators to get updates to software packages and track information about support contracts and subscriptions used by the systems. Standard tools such as **PackageKit** and **yum** can obtain software packages and updates through a content distribution network provided by Red Hat.

### Using Red Hat Subscription Management tools

There are four basic tasks performed with Red Hat Subscription Management tools:

- **Register** a system to associate that system to a Red Hat account. This allows Subscription Manager to uniquely inventory the system. When no longer in use, a system may be unregistered.
- **Subscribe** a system to entitle it to updates for selected Red Hat products. Subscriptions have specific levels of support, expiration dates, and default repositories. The tools can be used to either auto-attach or select a specific entitlement. As needs change, subscriptions may be removed.
- **Enable repositories** to provide software packages. Multiple repositories are enabled by default with each subscription, but other repositories such as updates or source code can be enabled or disabled as needed.
- **Review and track** entitlements which are available or consumed. Subscription information can be viewed locally on a specific system or, for an account, in either the Red Hat Customer Portal **Subscriptions** page or the Subscription Asset Manager (SAM).

### Register a system

To register a system with the subscription management service, launch **subscription-manager-gui** by selecting **Applications** → **Other** → **Red Hat Subscription Manager** from the main GNOME menu. Enter the password for **root** when prompted to authenticate. This will display the following **Subscription Manager** window.

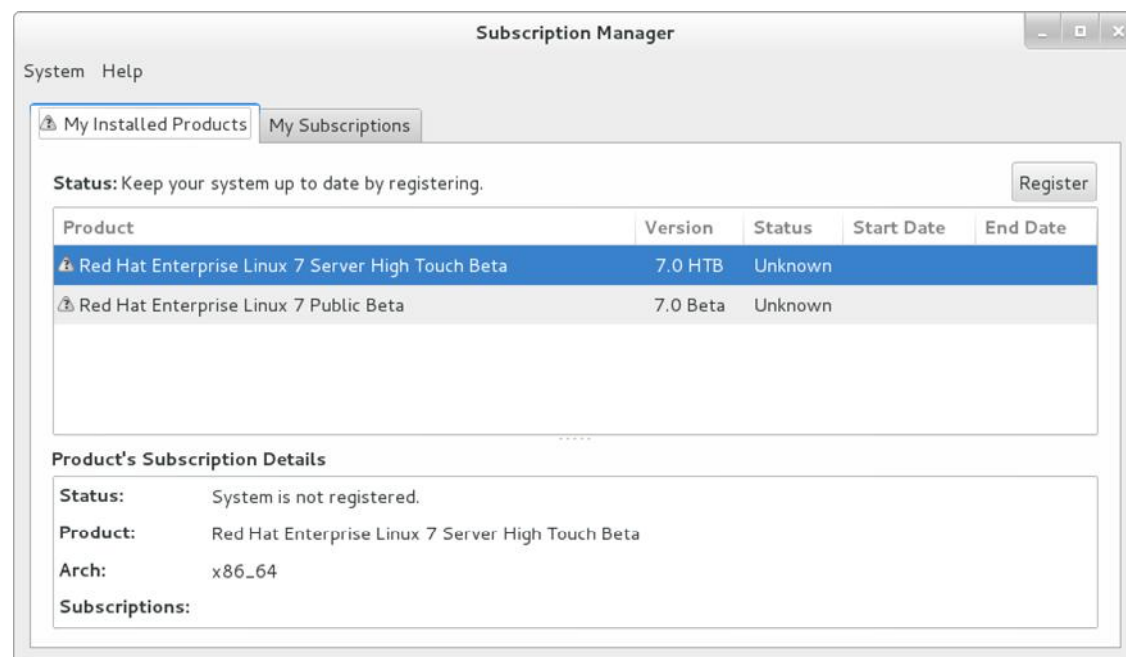
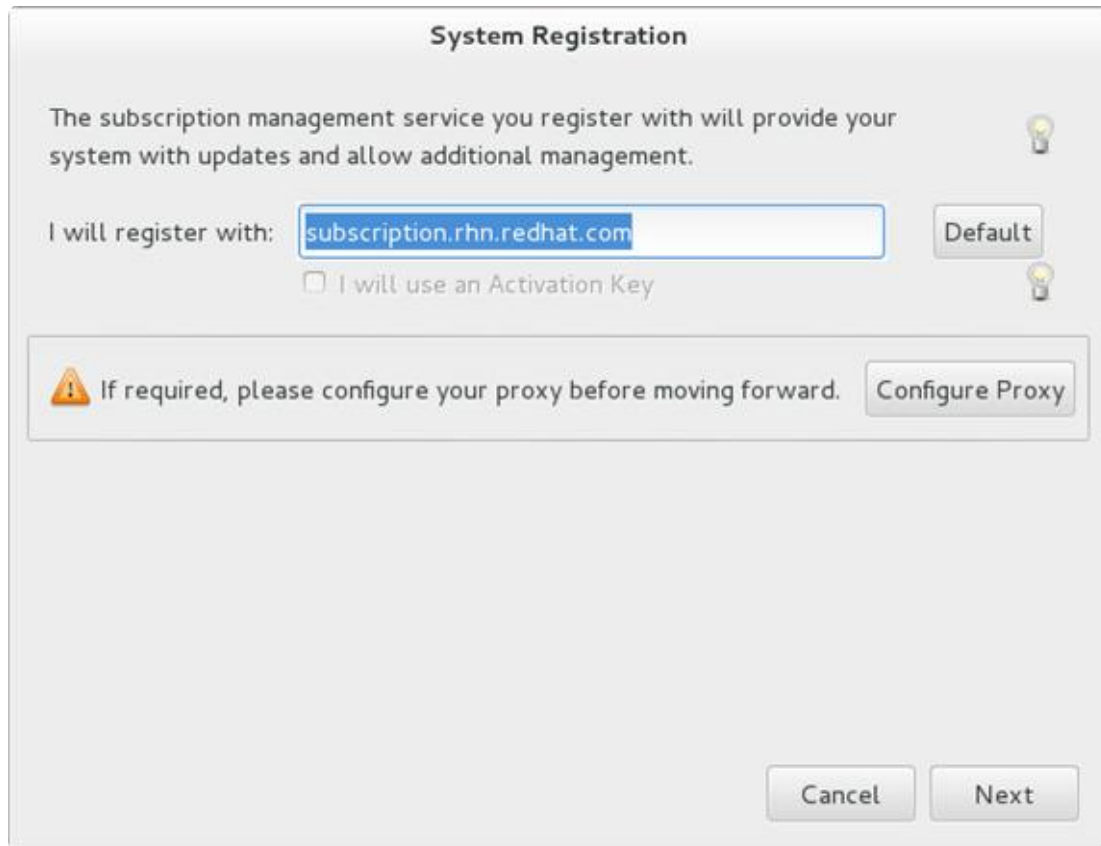



Figure 13.1: The main window of Red Hat Subscription Manager

To register the system, click the **Register** button in the top-right corner of the **Subscription Manager** window. This will bring up the following dialog:



*Figure 13.2: The service location dialog of Red Hat Subscription Manager*

This dialog box registers a system with a subscription server. The default (subscription.rhn.redhat.com) will register the server to Red Hat's "hosted" content distribution network.



The image shows a 'System Registration' dialog box with a light gray background. At the top, the title 'System Registration' is centered. Below the title, there are two sections. The first section is titled 'Please enter your Red Hat account information:' and contains two text input fields: 'Login:' with the text 'demo-account' and 'Password:' with seven black dots. Below these fields is a tip icon (a lightbulb) and the text: 'Tip: Forgot your login or password? Look it up at [http://redhat.com/forgot\\_password](http://redhat.com/forgot_password)'. The second section is titled 'Please enter the following for this system:' and contains one text input field labeled 'System Name:' with the text 'demo.example.com'. Below this field is a checked checkbox with the label 'Manually attach subscriptions after registration'. At the bottom right of the dialog box are two buttons: 'Cancel' and 'Register'.

*Figure 13.3: The account information dialog of Red Hat Subscription Manager*

Click **Next**, then authenticate using the Red Hat account to which the system should be registered. By default, **Subscription Manager** will try to find the best subscription for this system out of all available subscriptions. If more than one subscription is available, or a specific subscription is required, check the **Manually attach subscriptions after registration** checkbox. With this option checked, **Subscription Manager** will only register the system and not automatically assign any subscriptions.

Click the **Register** button to complete the registration.

#### **Assigning subscriptions**

To assign subscriptions to a system, navigate to the **All Available Subscriptions** tab in the main window of **Subscription Manager**, then click the **Update** button to retrieve a list of available subscriptions.

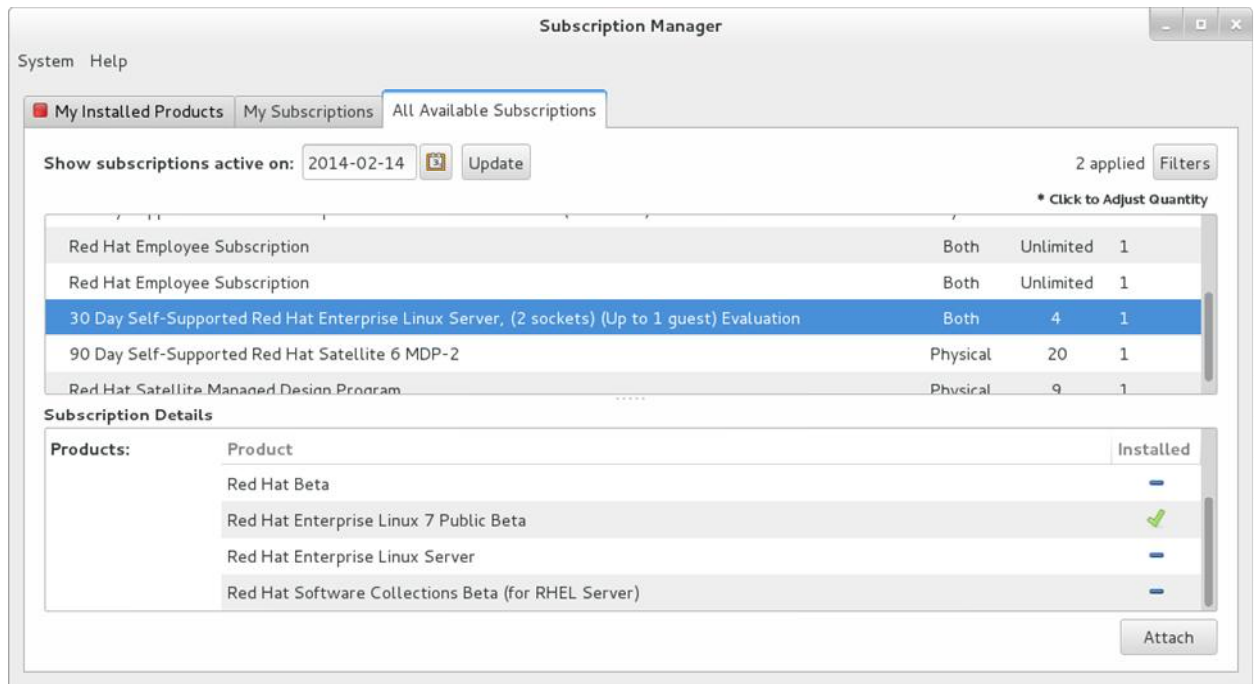


Figure 13.4: All Available Subscriptions tab of Red Hat Subscription Manager

From this list, select one or more subscriptions to assign to this system, then click the **Attach** button. If there is more than one contract for a specific subscription, a new dialog will open up, asking you to select which contract to use. Note that there are different contracts for *Physical* and *Virtual* systems.

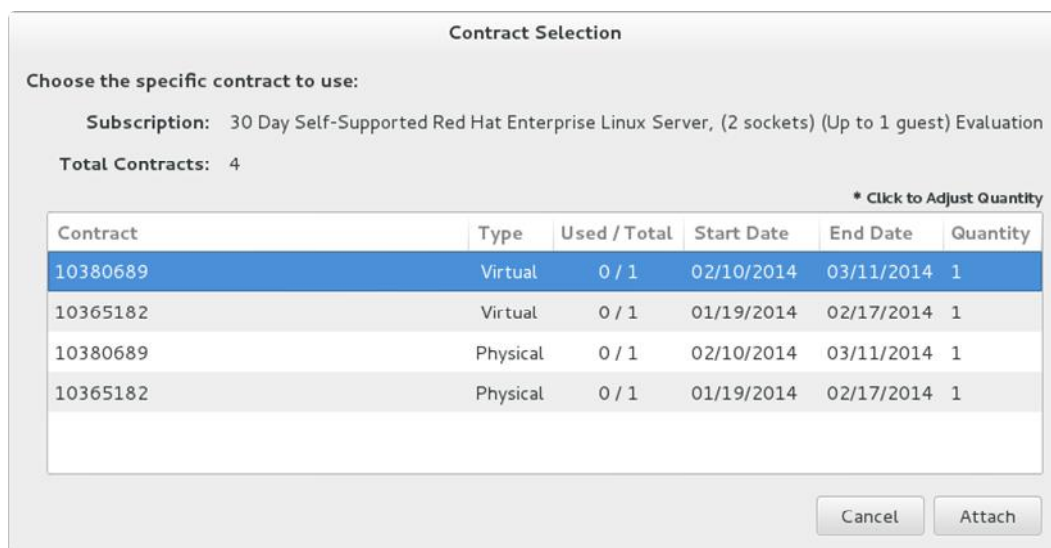


Figure 13.5: Contract selection dialog of Red Hat Subscription Manager

After a subscription has been assigned, close the **Subscription Manager** window. The system is now properly subscribed and ready to receive updates and/or install new software from Red Hat.

### Automating registrations and subscriptions

Use **subscription-manager(8)** to register a system without using a graphical environment. The **subscription-manager** command can automatically attach a system to the best-matched compatible subscriptions for the system.

- Register a system to a Red Hat account:



- `[root:@serverX ~]# subscription-manager register --username=yourusername --password=yourpassword`

- View available subscriptions:

- `[root:@serverX ~]# subscription-manager list --available | less`

- Auto-attach a subscription:

- `[root:@serverX ~]# subscription-manager attach --auto`

- View consumed subscriptions:

- `[root:@serverX ~]# subscription-manager list --consumed`

- Unregister a system:

- `[root:@serverX ~]# subscription-manager unregister`

#### Note

**subscription-manager** can also be used in conjunction with *activation keys*, allowing registration and assignment of predefined subscriptions, without using a username or password. This method of registration can be very useful for automated installations and deployments. Activation keys are usually issued by an on-premise subscription management service, such as Subscription Asset Manager, and will not be discussed in detail in this course.

#### Entitlement certificates

An entitlement is a subscription that's been attached to a system. Digital certificates are used to store current information about entitlements on the local system. Once registered, the entitlement certificates are stored in `/etc/pki` and its subdirectories.

- `/etc/pki/product` contains certificates which indicate Red Hat products installed on the system.
- `/etc/pki/consumer` contains certificates which indicate the Red Hat account to which the system is registered.
- `/etc/pki/entitlement` contains certificates which indicate which subscriptions are attached to the system.

The certificates can be inspected with the **rct** utility directly, but normally the **subscription-manager** tools are a more user-friendly way to examine the subscriptions that are attached to the system.

#### Important

Older versions of Red Hat Enterprise Linux originally supported a different subscription management method, *RHN Classic*. RHN Classic is not supported by Red Hat Enterprise Linux 7.

The method covered in this section, *Red Hat Subscription Management*, is the only one used by RHEL 7, and is the default method used by RHEL 6 after RHEL 6.3, and RHEL 5 after RHEL 5.9. RHEL 4 only

supports the old method. More information on both methods is available in the references at the end of this section.

## References

**subscription-manager-gui(8)**, **subscription-manager(8)**, and **rct(8)** man pages

[Get started with Red Hat Subscription Management](#)

[Red Hat Subscription Management: Migrating from RHN and Satellite](#)

## 13.2. Practice: Red Hat Subscription Management

### Quiz

Match the following items to their counterparts in the table.

| Description                                      | Task                |
|--------------------------------------------------|---------------------|
| Determine the number of available subscriptions  | Review and track    |
| Enable a system to use selected Red Hat products | Subscribe           |
| Attach a system to a Red Hat account             | Register            |
| Provide software packages                        | Enable repositories |

## 13.3. RPM Software Packages and Yum


### Objectives

After completing this section, students should be able to explain what an RPM package is and how RPM packages are used to manage software on a Red Hat Enterprise Linux system.

### Software packages and RPM

Many years ago, Red Hat developed the **RPM Package Manager** which provides a standard way to package software for distribution. Managing software in the form of *RPM packages* is much simpler than working with software that has simply been extracted into a file system from an archive. It lets administrators track which files were installed by the software package and which ones need to be removed if it is uninstalled, and check to ensure supporting packages are present when it is installed. Information about installed packages is stored in a local RPM database on each system. All software provided by Red Hat for Red Hat Enterprise Linux is provided as an RPM package.

RPM package files are named using a combination of the package **name-version-release.architecture**:



- NAME is one or more words describing the contents (httpd-tools).
- VERSION is the version number of the original software (2.4.6).
- RELEASE is the release number of the package based on that version, and is set by the packager, who might not be the original software developer (7.el7).
- ARCH is the processor architecture the package was compiled to run on. "noarch" indicates that this package's contents are not architecture-specific (x86\_64).

When installing packages from repositories, only the package name is required. The package with the higher version will be installed. If there are multiple files with the same version, the package with the higher release number will be installed.

Each RPM package is a special archive made up of three components:

- The files installed by the package.
- Information about the package (metadata), such as the name/version/release/arch; a summary and description of the package; whether it requires other packages to be installed; licensing; a package changelog; and other details.
- Scripts which may run when this package is installed, updated, or removed, or which are triggered when other packages are installed, updated, or removed.

RPM packages can be digitally signed by the organization that packaged them. All packages from a particular source are normally signed by the same GPG private key. If the package is altered or corrupted, the signature will no longer be valid. This allows the system to verify package integrity before installing them. All RPM packages released by Red Hat are digitally signed.

### Updates and Patches

When the upstream source code for a software package is patched by Red Hat, a complete RPM package is generated. If a package is newly added to a system, only the latest version of that package is needed, not every version of the package since the first release. For systems that need updating, the old version of the package is actually removed and the new version is installed. Configuration files are usually retained during an upgrade, but the exact behavior for a particular package is defined when the new version of the package is created.

In most cases, only one version or release of a package may be installed at a time. Typically, the RPM installation process will not allow files to be overwritten. If a package is built so that there are no conflicting filenames, then multiple versions may be installed. This is the case for the **kernel** package. Since a new kernel can only be tested by booting to that kernel, the package is specifically designed so that multiple versions may be installed at once. If the new kernel fails to boot, the old kernel is still available.

## ▣ The yum package manager

Once a system is installed, additional software packages and updates are normally installed from a network *package repository*, most frequently through the Red Hat subscription management service discussed in the previous section. The **rpm** command may be used to install, update, remove, and query RPM packages. However, it does not resolve dependencies automatically and all packages must be listed. Tools such as **PackageKit** and **yum** are front-end applications for **rpm** and can be used to install individual packages or *package collections* (sometimes called *package groups*).

The **yum** command searches numerous repositories for packages and their dependencies so they may be installed together in an effort to alleviate dependency issues. The main configuration file for **yum** is **/etc/yum.conf** with additional repository configuration files located in the **/etc/yum.repos.d** directory. Repository configuration files include, at a minimum, a repo id (in square brackets), a name and the URL location of the package repository. The URL can point to a local directory (file) or remote network share (http, ftp, etc.). If the URL is pasted in a browser, the contents should display the RPM packages, possibly in one or more subdirectories, and a **repodata** directory with information about available packages.

The **yum** command is used to list repositories, packages, and package groups:

```
[root@serverX ~]# yum repolist
Loaded plugins: langpacks
repo id repo name status
!rhel_dvd Remote classroom copy of dvd 4,529
repolist: 4,529

[root@serverX ~]# yum list yum*
Loaded plugins: langpacks
Installed Packages
yum.noarch 3.4.3-118.2.e17 @anaconda/7.0
yum-langpacks.noarch 0.4.2-3.e17 @rhel_dvd
yum-metadata-parser.x86_64 1.1.4-10.e17 @anaconda/7.0
yum-rhn-plugin.noarch 2.0.1-4.e17 @rhel_dvd
yum-utils.noarch 1.1.31-24.e17 @rhel_dvd
Available Packages
yum-plugin-aliases.noarch 1.1.31-24.e17 rhel_dvd
yum-plugin-changelog.noarch 1.1.31-24.e17 rhel_dvd
yum-plugin-tmprepo.noarch 1.1.31-24.e17 rhel_dvd
yum-plugin-verify.noarch 1.1.31-24.e17 rhel_dvd
yum-plugin-versionlock.noarch 1.1.31-24.e17 rhel_dvd

[root@serverX ~]# yum list installed
Loaded plugins: langpacks
Installed Packages
GConf2.x86_64 3.2.6-8.e17 @rhel_dvd
ModemManager.x86_64 1.1.0-6.git20130913.e17 @rhel_dvd
ModemManager-glib.x86_64 1.1.0-6.git20130913.e17 @rhel_dvd
...

[root@serverX ~]# yum grouplist
...
Installed groups:
```

```
Base
Desktop Debugging and Performance Tools
Dial-up Networking Support
Fonts
Input Methods
...
```

**References**

**yum(8)**, **yum.conf(5)**, **rpm(8)**, **rpm2cpio(8)**, and **rpmkeys(8)** man pages

### 13.4. Practice: RPM Software Packages

#### Quiz

Match the following items to their counterparts in the table.

| Description                                          | Term          |
|------------------------------------------------------|---------------|
| The upstream source code version                     | Version       |
| List of reasons for each package build               | Changelog     |
| The version of the package build                     | Release       |
| The processor type required for a specific package   | Architecture  |
| A collection of RPM packages and package groups      | Repository    |
| Used to verify the source and integrity of a package | GPG signature |

### 13.5. Managing Software Updates with yum

#### Objectives

After completing this section, students should be able to find, install, and update software packages using the **yum** command.

#### Working with yum

##### Using yum

**yum** is a powerful command-line tool that can be used to more flexibly manage (install, update, remove, and query) software packages. Official Red Hat packages are normally downloaded from Red Hat's content distribution network. Registering a system to the subscription management service automatically configures access to software repositories based on the attached subscriptions.

##### Finding software with yum

- o **yum help** will display usage information.
- o **yum list** displays installed and available packages.

```
[root@serverX ~]# yum list 'http*'
Loaded plugins: langpacks
```

|                                |              |          |
|--------------------------------|--------------|----------|
| ○ Available Packages           |              |          |
| ○ httpcomponents-client.noarch | 4.2.5-4.e17  | rhel_dvd |
| ○ httpcomponents-core.noarch   | 4.2.4-6.e17  | rhel_dvd |
| ○ httpd.x86_64                 | 2.4.6-17.e17 | rhel_dvd |
| ○ httpd-devel.x86_64           | 2.4.6-17.e17 | rhel_dvd |
| ○ httpd-manual.noarch          | 2.4.6-17.e17 | rhel_dvd |
| ○ httpd-tools.x86_64           | 2.4.6-17.e17 | rhel_dvd |

- **yum search *KEYWORD*** lists packages by keywords found in the name and summary fields only.

To search for packages that have "web server" in their name, summary, and description fields, use **search all**:

```
[root@serverX ~]# yum search all 'web server'
```

Loaded plugins: langpacks

```
===== Matched: web server =====
```

```
freeradius.x86_64 : High-performance and highly configurable free RADIUS server
hsqldb.noarch : HyperSQL Database Engine
httpd.x86_64 : Apache HTTP Server
libcurl.i686 : A library for getting files from web servers
libcurl.x86_64 : A library for getting files from web servers
mod_revocator.x86_64 : CRL retrieval module for the Apache HTTP server
mod_security.x86_64 : Security module for the Apache HTTP Server
python-paste.noarch : Tools for using a Web Server Gateway Interface stack
```

- **yum info *PACKAGENAME*** gives detailed information about a package, including the disk space needed for installation.

To get information on the Apache HTTP Server:

```
[root@serverX ~]# yum info httpd
```

Loaded plugins: langpacks

```
Available Packages
```

```
Name : httpd
Arch : x86_64
Version : 2.4.6
Release : 17.e17
Size : 1.1 M
Repo : rhel_dvd
Summary : Apache HTTP Server
```

```
URL : http://httpd.apache.org/
License : ASL 2.0
Description : The Apache HTTP Server is a powerful, efficient, and extensible
 : web server.
```

- **yum provides *PATHNAME*** displays packages that match the pathname specified (which often include wildcard characters).

To find packages that provide the `/var/www/html` directory, use:

```
[root@serverX ~]# yum provides /var/www/html
Loaded plugins: langpacks
httpd-2.4.6-17.el7.x86_64 : Apache HTTP Server
Repo : rhel_dvd
Matched from:
Filename : /var/www/html

1:php-pear-1.9.4-21.el7.noarch : PHP Extension and Application Repository
 : framework
Repo : rhel_dvd
Matched from:
Filename : /var/www/html
```

### Installing and removing software with yum

- **yum install *PACKAGENAME*** obtains and installs a software package, including any dependencies.

```
○ [root@serverX ~]# yum install httpd
○ Loaded plugins: langpacks
○ Resolving Dependencies
○ --> Running transaction check
○ ---> Package httpd.x86_64 0:2.4.6-17.el7 will be installed
○ --> Processing Dependency: httpd-tools = 2.4.6-17.el7 for package:
httpd-2.4.6-17.el7.x86_64
○ --> Processing Dependency: /etc/mime.types for package:
httpd-2.4.6-17.el7.x86_64
○ --> Processing Dependency: libapr-1.so.0() (64bit) for package:
httpd-2.4.6-17.el7.x86_64
```

```

○ --> Processing Dependency: libaprutil-1.so.0() (64bit) for package:
httpd-2.4.6-17.el7.x86_64

○ --> Running transaction check

○ ---> Package apr.x86_64 0:1.4.8-3.el7 will be installed

○ ---> Package apr-util.x86_64 0:1.5.2-6.el7 will be installed

○ ---> Package httpd-tools.x86_64 0:2.4.6-17.el7 will be installed

○ ---> Package mailcap.noarch 0:2.1.41-2.el7 will be installed

○ --> Finished Dependency Resolution

○

○ Dependencies Resolved

○

○ =====
====
○ Package Arch Version Repository Size
○ =====
====
○ Installing:
○ httpd x86_64 2.4.6-17.el7 rhel_dvd 1.1 M
○ Installing for dependencies:
○ apr x86_64 1.4.8-3.el7 rhel_dvd 100 k
○ apr-util x86_64 1.5.2-6.el7 rhel_dvd 90 k
○ httpd-tools x86_64 2.4.6-17.el7 rhel_dvd 76 k
○ mailcap noarch 2.1.41-2.el7 rhel_dvd 31 k
○

○ Transaction Summary
○ =====
====
○ Install 1 Package (+4 Dependent packages)
○

○ Total download size: 1.4 M
○ Installed size: 4.3 M
○ Is this ok [y/d/N]:

```

- **yum update *PACKAGENAME*** obtains and installs a newer version of the software package, including any dependencies. Generally the process tries to preserve configuration files in place, but in some cases, they may be renamed if the packager thinks the old one will not work after the update. With no *PACKAGENAME* specified, it will install all relevant updates.



```
[root@serverX ~]# yum update
```

Since a new kernel can only be tested by booting to that kernel, the package is specifically designed so that multiple versions may be installed at once. If the new kernel fails to boot, the old kernel is still available. Using `yum update kernel` will actually *install* the new kernel. The configuration files hold a list of packages to "always install" even if the administrator requests an update.

#### Note

Use `yum list kernel` to list all installed and available kernels. To view the currently running kernel, use the `uname` command. The `-r` option will show only the kernel version and release, and the `-a` option will show the kernel release and additional information.

```
[root@serverX ~]# yum list kernel
```

```
Loaded plugins: langpacks
```

```
Installed Packages
```

```
kernel.x86_64 3.10.0-123.0.1.el7 @anaconda/7.0
kernel.x86_64 3.10.0-84.el7 @rhel-7-server-htb-rpms
```

```
[root@serverX ~]# uname -r
```

```
3.10.0-123.el7.x86_64
```

```
[root@serverX ~]# uname -a
```

```
Linux demo.example.com 3.10.0-123.el7.x86_64 #1 SMP Tue Nov 26 16:51:22 EST 2013
x86_64 x86_64 x86_64 GNU/Linux
```

- `yum remove PACKAGENAME` removes an installed software package, including any supported packages.

```
[root@serverX ~]# yum remove httpd
```

#### Warning

`yum remove` will remove the package(s) listed *and any package that requires the package(s) being removed*(and package(s) which require those packages, and so on). This can lead to unexpected removal of packages, so carefully check the list of packages to be removed.

### Installing and removing groups of software with `yum`

- `yum` also has the concept of *groups*, which are collections of related software installed together for a particular purpose. In Red Hat Enterprise Linux 7, there are two kinds of groups. Regular groups are collections of packages. *Environment groups* are collections of other groups which include their own packages. The packages or groups provided by a group may be *mandatory* (must be installed if the group is installed), *default* (are normally installed if the group is installed), or *optional* (are not installed when the group is unless asked for specifically).

Like **yum list**, the **yum group list** (or **yum grouplist**) command will show the names of installed and available groups. Some groups are normally installed through environment groups and are hidden by default. These hidden groups can also be listed with the **yum group list hidden** command. If the **ids** option is added, the group ID will also be shown. Groups can be installed, updated, removed, and otherwise queried by name or ID.

```
[root@serverX ~]# yum group list
Loaded plugins: langpacks
Available environment groups:
 Minimal install
 Infrastructure Server
 File and Print Server
 Web Server
 Virtualization Host
 Server with GUI
Installed groups:
 Base
 Desktop Debugging and Performance Tools
 Dial-up Networking Support
 Fonts
 Input Methods
 Internet Browser
 PostgreSQL Database server
 Printing client
 X Window System
Available Groups:
 Additional Development
 Backup Client
 Backup Server
 ...
```

- Information about a group is displayed with **yum group info** (or **yum groupinfo**). It includes a list of mandatory, default, and optional package names or group IDs. The package names or group IDs may have a marker in front of them.

| Marker | Meaning                                                               |
|--------|-----------------------------------------------------------------------|
| =      | Package is installed, was installed as part of the group              |
| +      | Package isn't installed, will be if the group is installed or updated |

| Marker    | Meaning                                                                   |
|-----------|---------------------------------------------------------------------------|
| -         | Package isn't installed, will not be if the group is installed or updated |
| no marker | Package is installed, but was not installed through the group.            |

```

o [root@serverX ~]# yum group info "Identity Management
 Server"
o Loaded plugins: langpacks
o
o Group: Identity Management Server
o Group-Id: identity-management-server
o Description: Centralized management of users, servers and authentication
 policies.
o Default Packages:
o +389-ds-base
o +ipa-admintools
o +ipa-server
o +pki-ca
o Optional Packages:
o +ipa-server-trust-ad
o +nuxwdog
o +slapi-nis

```

- o The **yum group install** (or **yum groupinstall**) command will install a group which will install its mandatory and default packages and the packages they depend on.

```

o [root@serverX ~]# yum group install "Infiniband
 Support"
o ...
o Transaction Summary
o =====
o ====
o Install 17 Packages (+7 Dependent packages)
o
o Total download size: 9.0 M
o Installed size: 33 M
o Is this ok [y/d/N]:

```

```

o ...

```

**Important**

The behavior of **yum** groups has changed in Red Hat Enterprise Linux 7 from Red Hat Enterprise Linux 6 and earlier. In RHEL 7, groups are treated as *objects*, and are tracked by the system. If an installed group is updated, and new mandatory or default packages have been added to the group by the **yum** repository, those new packages will be installed on update.

RHEL 6 and earlier consider a group to be installed if all its mandatory packages have been installed; or if it had no mandatory packages, if any default or optional packages in the group are installed. In RHEL 7, a group is considered to be installed *only* if **yum group install** was used to install it. A new command in RHEL 7, **yum group mark install GROUPNAME** can be used to mark a group as installed, and any missing packages and their dependencies will be installed on the next update.

Finally, RHEL 6 and earlier did not have the two-word form of the **yum group** commands. In other words, in RHEL 6 the command **yum grouplist** existed, but the equivalent RHEL 7 command **yum group list** did not.

### Viewing transaction history

- All install and **remove tra**nsactions are logged in **/var/log/yum.log**.

```
[root@serverX ~]# tail -5 /var/log/yum.log
Feb 16 14:10:41 Installed: libnes-1.1.3-5.e17.x86_64
Feb 16 14:10:42 Installed: libmthca-1.0.6-10.e17.x86_64
Feb 16 14:10:43 Installed: libmlx4-1.0.5-7.e17.x86_64
Feb 16 14:10:43 Installed: libibcm-1.0.5-8.e17.x86_64
Feb 16 14:10:45 Installed: rdma-7.0_3.13_rc8-3.e17.noarch
```

- A summary of install and remove transactions can be viewed with **yum history**.

```
[root@serverX ~]# yum history
Loaded plugins: langpacks
ID	Login user	Date and time	Action(s)	Altered

 6 | Student User <student> | 2014-02-16 14:09 | Install | 25
 5 | Student User <student> | 2014-02-16 14:01 | Install | 1
 4 | System <unset> | 2014-02-08 22:33 | Install | 1112 EE
 3 | System <unset> | 2013-12-16 13:13 | Erase | 4
 2 | System <unset> | 2013-12-16 13:13 | Erase | 1
 1 | System <unset> | 2013-12-16 13:08 | Install | 266
history list
```

- A transaction can be reversed with the **history undo** options:

```
[root@serverX ~]# yum history undo 6
```

```

○ Loaded plugins: langpacks
○ Undoing transaction 6, from Sun Feb 16 14:09:51 2014
○ Install dap1-2.0.39-2.el7.x86_64 @rhel-7-server-htb-rpms
○ Dep-Install graphviz-2.30.1-18.el7.x86_64 @rhel-7-server-htb-rpms
○ Dep-Install graphviz-tcl-2.30.1-18.el7.x86_64 @rhel-7-server-htb-rpms
○ Install ibacm-1.0.8-4.el7.x86_64 @rhel-7-server-htb-rpms
○ Install ibutils-1.5.7-9.el7.x86_64 @rhel-7-server-htb-rpms
○ Dep-Install ibutils-libs-1.5.7-9.el7.x86_64 @rhel-7-server-htb-rpms
...

```

## Summary of yum commands

Packages can be located, installed, updated, and removed by name or by package groups.

| Task:                                         | Command:                                  |
|-----------------------------------------------|-------------------------------------------|
| List installed and available packages by name | <code>yum list [NAME-PATTERN]</code>      |
| List installed and available groups           | <code>yum grouplist</code>                |
| Search for a package by keyword               | <code>yum search KEYWORD</code>           |
| Show details of a package                     | <code>yum info PACKAGENAME</code>         |
| Install a package                             | <code>yum install PACKAGENAME</code>      |
| Install a package group                       | <code>yum groupinstall "GROUPNAME"</code> |
| Update all packages                           | <code>yum update</code>                   |
| Remove a package                              | <code>yum remove PACKAGENAME</code>       |
| Display transaction history                   | <code>yum history</code>                  |

**References**  
**yum**(1) and **yum.conf**(5) man pages  
 Additional information on **yum** may be available in the *Red Hat Enterprise Linux System Administrator's Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

## 13.6. Practice: Installing and Updating Software with yum

### Guided exercise

In this lab, you will install and remove packages and package groups.

**Outcomes:**

Explore installing and removing packages with dependencies.

**Before you begin...**

Reset your serverX system.

1. Search for a specific package.
  - a. Attempt to run the command **gnuplot**. You should find that it is not installed.

```
b. [root@serverX ~]# gnuplot
c. bash: gnuplot: command not found...
```

d. Search for plotting packages.

```
e. [root@serverX ~]# yum search plot
f. Loaded plugins: langpacks
g. ===== N/S matched: plot =====
h. emacs-gnuplot.noarch : Emacs bindings for the gnuplot main application
i. gnuplot.x86_64 : A program for plotting mathematical expressions and
 data
j. gnuplot-common.x86_64 : The common gnuplot parts
k. python-matplotlib.x86_64 : Python 2D plotting library
l. texlive-pst-plot.noarch : Plot data using PSTricks
m.
n. Name and summary matches only, use "search all" for everything.
```

o. Find out more information about the **gnuplot** package.

```
p. [root@serverX ~]# yum info gnuplot
q. Name : gnuplot
r. Arch : x86_64
```

...

2. Install the **gnuplot** package.

```
3. [root@serverX ~]# yum install -y gnuplot
4. ...
5. Dependencies Resolved
6.
7. =====
 =
8. Package Arch Version Repository Size
9. =====
 =
10. Installing:
11. gnuplot x86_64 4.6.2-3.e17 rhel_dvd 645 k
12. Installing for dependencies:
13. gnuplot-common x86_64 4.6.2-3.e17 rhel_dvd 595 k
```

```
14.
15. Transaction Summary
16. =====
 =
17. Install 1 Package (+1 Dependent package)
```

```
...
```

18. Remove packages.

- a. Attempt to remove the **gnuplot** package, but say no. How many packages would be removed?

```
b. [root@serverX ~]# yum remove gnuplot
c. ...
d. Removing:
e. gnuplot x86_64 4.6.2-3.e17 @rhel_dvd 1.5 M
f.
g. Transaction Summary
h. =====
 =====
i. Remove 1 Package
j.
k. Installed size: 1.5 M
```

```
Is this ok [y/N]: n
```

- l. Attempt to remove the **gnuplot-common** package, but say no. How many packages would be removed?

```
m. [root@serverX ~]# yum remove gnuplot-common
n. ...
o. Removing:
p. gnuplot-common x86_64 4.6.2-3.e17 @rhel_dvd 1.4 M
q. Removing for dependencies:
r. gnuplot x86_64 4.6.2-3.e17 @rhel_dvd 1.5 M
s.
t. Transaction Summary
u. =====
 =====
v. Remove 1 Package (+1 Dependent package)
```

```
w.
x. Installed size: 2.9 M
y.
```

```
Is this ok [y/N]: n
```

19. Gather information about the "Compatibility Libraries" component group and install it on serverX.

a. List all available component groups.

```
[root@serverX ~]# yum grouplist
```

b. Find out more information about the *Compatibility Libraries* component group, including a list of included packages.

```
c. [root@serverX ~]# yum groupinfo "Compatibility
Libraries"
d. Loaded plugins: langpacks
e.
f. Group: Compatibility Libraries
g. Group-Id: compat-libraries
h. Description: Compatibility libraries for applications built on
previous versions of Red Hat Enterprise Linux.
i. Mandatory Packages:
j. +compat-db47
k. +compat-glibc
l. +compat-libcap1
m. +compat-libf2c-34
n. +compat-libgfortran-41
o. +compat-libtiff3
p. +compat-openldap
q. +libpng12
r. +openssl1098e
```

s. Install the *Compatibility Libraries* component group.

```
t. [root@serverX ~]# yum groupinstall
"Compatibility Libraries"
u. Loaded plugins: langpacks
v. Resolving Dependencies
w. --> Running transaction check
```



```

x. --> Package compat-db47.x86_64 0:4.7.25-27.e17 will be installed
y. --> Processing Dependency: compat-db-headers = 4.7.25-27.e17 for
package: compat-db47-4.7.25-27.e17.x86_64
z. ...
aa. Dependencies Resolved
bb.
cc. =====
 ===
dd. Package Arch Version Repository
ee. =====
 ===
ff. Installing for group install "Compatibility Libraries":
gg. compat-db47 x86_64 4.7.25-27.e17 rhel_dvd
hh. libpng12 x86_64 1.2.50-6.e17 rhel_dvd
ii. ...
jj. Installing for dependencies:
kk. compat-db-headers noarch 4.7.25-27.e17 rhel_dvd
ll. ...
mm.
nn. Transaction Summary
oo. =====
 ===
pp. Install 9 Packages (+3 Dependent packages)
qq.
rr. Total download size: 5.5 M
ss. Installed size: 21 M
tt. Is this ok [y/d/N]: y
uu. ...
vv. Installed:
ww. compat-db47.x86_64 0:4.7.25-27.e17
xx. compat-glibc.x86_64 1:2.12-4.e17
yy. ...
zz.
aaa. Dependency Installed:
bbb. compat-db-headers.noarch 0:4.7.25-27.e17
ccc. compat-glibc-headers.x86_64 1:2.12-4.e17
ddd.

```

```
eee. Complete!
```

20. Explore the history and undo options of **yum**.

a. Display recent **yum** history.

```
b. [root@serverX ~]# yum history
c. Loaded plugins: langpacks
d. ID | Login user | Date and time | Action(s) | Altered
e. -----
f. 3 | root <root> | 2014-06-05 09:33 | Install | 12
g. 2 | root <root> | 2014-06-05 09:30 | Install | 2
h. 1 | System <unset> | 2014-06-02 20:27 | Install | 1112 EE
i. history list
```

j. Confirm that the last transaction is the group installation.

```
k. [root@serverX ~]# yum history info 3
l. Loaded plugins: langpacks
m. Transaction ID : 3
n. Begin time : Thu Jun 5 09:33:19 2014
o. Begin rpmdb : 1210:7c6b529424621773d5fe147315a53d558f726814
p. End time : 09:33:40 2014 (21 seconds)
q. End rpmdb : 1222:c283bc776b18b9578b87cdec68853f49b31ca0cc
r. User : root <root>
s. Return-Code : Success
t. Command Line : groupinstall Compatibility Libraries
u. Transaction performed with:
v. Installed rpm-4.11.1-16.el7.x86_64 installed
w. Installed yum-3.4.3-117.el7.noarch installed
x. Packages Altered:
y. Dep-Install compat-db-headers-4.7.25-27.el7.noarch @rhel_dvd
z. Install compat-db47-4.7.25-27.el7.x86_64 @rhel_dvd
aa. ...
```

```
history info
```

bb. Use undo options to remove the last set of packages installed.

```
[root@serverX ~]# yum history undo 3
```

## 13.7. Enabling yum Software Repositories

### Objectives

After completing this section, students should be able to enable and disable the use of Red Hat or third-party yum repositories.

### Enabling Red Hat software repositories

Registering a system to the subscription management service automatically configures access to software repositories based on the attached subscriptions. To view all available repositories:

```
[root@serverX ~]# yum repolist all
Loaded plugins: langpacks
repo id repo name status
rhel-7-public-beta-debug-rpms Red Hat Enterprise Linux 7 Public disabled
rhel-7-public-beta-rpms Red Hat Enterprise Linux 7 Public enabled: 8,520
rhel-7-public-beta-source-rpms Red Hat Enterprise Linux 7 Public disabled
repolist: 8,520
```

Enable and disable repositories with **yum-config-manager**. This will change the **enabled** parameter in the **/etc/yum.repos.d/redhat.repo** file.

```
[root@serverX ~]# yum-config-manager --enable
rhel-7-public-beta-debug-rpms
Loaded plugins: langpacks
===== repo: rhel-7-public-beta-debug-rpms
=====
[rhel-7-public-beta-debug-rpms]
async = True
bandwidth = 0
base_persistdir = /var/lib/yum/repos/x86_64/7Server
baseurl = https://cdn.redhat.com/content/beta/rhel/everything/7/x86_64/debug
cache = 0
cachedir = /var/cache/yum/x86_64/7Server/rhel-7-public-beta-debug-rpms
cost = 1000
deltarpm_percentage =
enabled = 1
...
```

### Enabling third-party software repositories

*Configuring yum repositories*

Third-party repositories are directories of software package files provided by a non-Red Hat source, which can be accessed by **yum** from a website, FTP server, or local file system. Yum repositories are used by non-Red Hat distributors of software, or for small collections of local packages. (For example, Adobe provides some of its free software for Linux through a yum repository.)

The **content.example.com** classroom server actually hosts yum repositories for this class. Put a file in the **/etc/yum.repos.d/** directory to enable support for a new third-party repository. Repository configuration files must end in **.repo**. The repository definition contains the URL of the repository, a name, whether to use GPG to check the package signatures, and if so, the URL pointing to the trusted GPG key.

### Using yum-config-manager

If the URL for a yum repository is known, a configuration file can be created with **yum-config-manager**.

```
[root@serverX ~]# yum-config-manager
--add-repo="http://dl.fedoraproject.org/pub/epel/beta/7/x86_64/"
86_64/"
Loaded plugins: langpacks
adding repo from: http://dl.fedoraproject.org/pub/epel/beta/7/x86_64/

[dl.fedoraproject.org_pub_epel_beta_7_x86_64_]
name=added from: http://dl.fedoraproject.org/pub/epel/beta/7/x86_64/
baseurl=http://dl.fedoraproject.org/pub/epel/beta/7/x86_64/
enabled=1
```

A file was created in the **/etc/yum.repos.d** directory with the output shown. This file can now be modified to provide a customized name and the location of the GPG key. Administrators should download the key to a local file rather than allowing **yum** to retrieve the key from an external source.

```
[EPEL]
name=EPEL 7 Beta
baseurl=http://dl.fedoraproject.org/pub/epel/beta/7/x86_64/
enabled=1
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-EPEL-7
```

### RPM configuration package for the repository

Some repositories provide this configuration file and GPG public key as part of an RPM package that can be downloaded and installed using **yum localinstall**. One example of this is the volunteer project EPEL (Extra Packages for Enterprise Linux), which provides software not supported by Red Hat but compatible with Red Hat Enterprise Linux.

Installing the Red Hat Enterprise Linux 7 EPEL repo package:

```
[root@serverX ~]# rpm --import
http://dl.fedoraproject.org/pub/epel/RPM-GPG-KEY-EPEL-7
```

```
[root@serverX ~]# yum install
http://dl.fedoraproject.org/pub/epel/beta/7/x86_64/epel-r
elease-7-0.1.noarch.rpm
```

Configuration files often list multiple repository references in a single file. Each repository reference begins with a single-word name in square brackets.

```
[root@serverX ~]# cat /etc/yum.repos.d/epel.repo

[epel]
name=Extra Packages for Enterprise Linux 7 - $basearch
#baseurl=http://download.fedoraproject.org/pub/epel/7/$basearch
mirrorlist=https://mirrors.fedoraproject.org/metalink?repo=epel-7&arch=$base
arch
failovermethod=priority
enabled=1
gpgcheck=0
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-EPEL-7

[epel-debuginfo]
name=Extra Packages for Enterprise Linux 7 - $basearch - Debug
#baseurl=http://download.fedoraproject.org/pub/epel/7/$basearch/debug
mirrorlist=https://mirrors.fedoraproject.org/metalink?repo=epel-debug-7&arch
=$basearch
failovermethod=priority
enabled=0
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-EPEL-7
gpgcheck=1

[epel-source]
name=Extra Packages for Enterprise Linux 7 - $basearch - Source
#baseurl=http://download.fedoraproject.org/pub/epel/7/SRPMS
mirrorlist=https://mirrors.fedoraproject.org/metalink?repo=epel-source-7&arc
h=$basearch
failovermethod=priority
enabled=0
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-EPEL-7
gpgcheck=1
```

The **enabled=0** parameter can be included so that a repository is defined but not searched by default. Repositories can be enabled and disabled persistently with **yum-config-manager** or temporarily with **--enablerepo=PATTERN** and **--disablerepo=PATTERN** options on **yum**.

#### Warning

Install the RPM GPG key before installing signed packages. This will verify that the packages belong to a key which has been imported. Otherwise, **yum** will complain about the missing key.

(The **--nogpgcheck** option can be used to ignore missing GPG keys, but this could cause forged or insecure packages to be installed on the system, potentially compromising its security.)

#### References

Additional information may be available in the section on configuring yum and yum repositories in the *Red Hat Enterprise Linux System Administrator's Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

**yum**(1), **yum.conf**(5), and **yum-config-manager**(1) man pages

## 13.8. Practice: Enabling Software Repositories

### Guided exercise

In this lab, you will configure your server to use a separate **yum** repository to obtain updates, and update your machine.

#### Outcomes:

The system will be configured to obtain software updates from a classroom server and the system will be running the latest Linux kernel.

#### Before you begin...

Reset your serverX system.

1. Configure the system to obtain software from two classroom repositories:
  - o Classroom packages provided at [http://content.example.com/rhel7.0/x86\\_64/rht](http://content.example.com/rhel7.0/x86_64/rht)
  - o Updates provided at [http://content.example.com/rhel7.0/x86\\_64/errata](http://content.example.com/rhel7.0/x86_64/errata)
- c. Use **yum-config-manager** to add the classroom packages repository.

```
d. [root@serverX ~]# yum-config-manager
 --add-repo="http://content.example.com/rhel7.0/
 x86_64/rht"
e. Loaded plugins: langpacks
f. adding repo from: http://content.example.com/rhel7.0/x86_64/rht
g.
h. [content.example.com_rhel7.0_x86_64_rht]
i. name=added from: http://content.example.com/rhel7.0/x86_64/rht
j. baseurl=http://content.example.com/rhel7.0/x86_64/rht
k. enabled=1
```

- l. Create the file `/etc/yum.repos.d/errata.repo` to enable the “Updates” repository with the following content:

```
m. [updates]
n. name=Red Hat Updates
o. baseurl = http://content.example.com/rhel7.0/x86_64/errata
p. enabled=1
q. gpgcheck=0
```

2. Use `yum-config-manager` to disable the classroom packages repository.

```
3. [root@serverX ~]# yum-config-manager --disable
 content.example.com_rhel7.0_x86_64_rht
4. Loaded plugins: langpacks
5. ===== repo: content.example.com_rhel7.0_x86_64_rht
 =====
6. [content.example.com_rhel7.0_x86_64_rht]
7. ...
8. enabled = 0
```

```
...
```

9. Update all relevant software provided using `yum update`.

```
[root@serverX ~]# yum update -y
```

10. Verify that there are two versions of the kernel installed. Which version is currently in use?

```
11. [root@serverX ~]# yum list kernel
12.
13. [root@serverX ~]# uname -r
14.
```

15. Reboot your serverX, then repeat the previous step. Which version is currently in use?

```
16. [root@serverX ~]# yum list kernel
17.
18. [root@serverX ~]# uname -r
19.
```

20. List, then install, the `rht-system` package.

```
21. [root@serverX ~]# yum list rht*
22. Loaded plugins: langpacks
```

```

23. Available Packages
24. rht-system.noarch 1.0.0-2.e17 updates
25. [root@serverX ~]# yum -y install rht-system
26. Loaded plugins: langpacks
27. Resolving Dependencies

```

```
...
```

## 13.9. Examining RPM Package Files

### Objectives

After completing this section, students should be able to examine and install downloaded package files.

### Examining downloaded packages with rpm

#### Examining RPM package files

The **rpm** utility is a low-level tool that can get information about the contents of package files and installed packages. It gets its information from a local database or the package files themselves.

The general form of a query is:

- **rpm -q [select-options] [query-options]**
- **rpm --query [select-options] [query-options]**

RPM queries: Select options

- **-q -a**: all installed packages
- **-q *PACKAGENAME***: currently installed *PACKAGENAME*

```

○ [root@serverX ~]# rpm -q yum
○ yum-3.4.3-118.e17.noarch

```

- **-q -p *PACKAGEFILE.rpm***: package file named *PACKAGEFILE.rpm*

```

○ [root@serverX ~]# rpm -q -p
 http://content.example.com/rhel7.0/x86_64/dvd/Packages/yum-utils-1.1.31-24.e17.noarch.rpm
○ yum-utils-1.1.31-24.e17.noarch.rpm

```

- **-q -f *FILENAME***: what package provides *FILENAME*

```

○ [root@serverX ~]# rpm -q -f /etc/yum/repos.d
○ yum-3.4.3-118.e17.noarch

```

RPM queries: Information about content of packages

- **-q**: lists the package's name and version; compare to **yum list**



- **-q -i**: package information; compare to **yum info**
- **-q -l**: list of files installed by the specified package

```
[root@serverX ~]# rpm -q -l yum-rhn-plugin
/etc/yum/pluginconf.d/rhnplugin.conf
/usr/share/doc/yum-rhn-plugin-2.0.1
/usr/share/doc/yum-rhn-plugin-2.0.1/LICENSE
/usr/share/locale/aln/LC_MESSAGES/yum-rhn-plugin.mo
```

...

- **-q -c**: list just the configuration files

```
[root@serverX ~]# rpm -q -c yum-rhn-plugin
/etc/yum/pluginconf.d/rhnplugin.conf
```

- **-q -d**: list just the documentation files

```
[root@serverX ~]# rpm -q -d yum-rhn-plugin
/usr/share/doc/yum-rhn-plugin-2.0.1/LICENSE
/usr/share/man/man5/rhnplugin.conf.5.gz
/usr/share/man/man8/rhnplugin.8.gz
/usr/share/man/man8/yum-rhn-plugin.8.gz
```

- **-q --scripts**: list shell scripts that may run before or after the package is installed or removed

```
[root@serverX ~]# rpm -q --scripts openssh-server
preinstall scriptlet (using /bin/sh):
getent group sshd >/dev/null || groupadd -g 74 -r sshd || :
getent passwd sshd >/dev/null || \
 useradd -c "Privilege-separated SSH" -u 74 -g sshd \
 -s /sbin/nologin -r -d /var/empty/sshdsd sshd 2> /dev/null || :
postinstall scriptlet (using /bin/sh):
○
○
if [$1 -eq 1] ; then
 # Initial installation
 /usr/bin/systemctl preset sshd.service sshd.socket >/dev/null 2>&1 || :
fi
preuninstall scriptlet (using /bin/sh):
○
```

```

o if [$1 -eq 0] ; then
o # Package removal, not upgrade
o /usr/bin/systemctl --no-reload disable sshd.service sshd.socket >
/dev/null 2>&1 || :
o /usr/bin/systemctl stop sshd.service sshd.socket > /dev/null 2>&1 || :
o fi
o postuninstall scriptlet (using /bin/sh):
o
o /usr/bin/systemctl daemon-reload >/dev/null 2>&1 || :
o if [$1 -ge 1] ; then
o # Package upgrade, not uninstall
o /usr/bin/systemctl try-restart sshd.service >/dev/null 2>&1 || :
o fi

```

- o **-q --changelog**: list change information for the package

```

o [root@serverX ~]# rpm -q --changelog audit
o * Thu Oct 03 2013 Steve Grubb <sgrubb@redhat.com> 2.3.2-3
o resolves: #828495 - semanage port should generate an audit event
o
o * Thu Aug 29 2013 Steve Grubb <sgrubb@redhat.com> 2.3.2-2
o resolves: #991056 - ausearch ignores USER events with -ua option

```

...

Querying and installing package files:

```

[root@serverX ~]# wget
http://classroom/pub/materials/wonderwidgets-1.0-4.x86_64
.rpm
[root@serverX ~]# rpm -q -l -p wonderwidgets-1.0-4.x86_64.rpm
/etc/wonderwidgets.conf
/usr/bin/wonderwidgets
/usr/share/doc/wonderwidgets-1.0
/usr/share/doc/wonderwidgets-1.0/README.txt

```

#### Note

The **repoquery** command can also be used to get information about packages and their contents. It differs from **rpm** by looking up that information in yum's repositories instead of the local database of installed packages.

## Using yum to install local package files

`yum localinstall PACKAGEFILE.rpm` can be used to install package files directly. It automatically downloads any dependencies the package has from any configured `yum` repositories. Packages are normally digitally signed to ensure they are legitimate; if the package is not signed by a key trusted by your system, it will be rejected. The `--nogpgcheck` option can disable the signature check if you are certain the package is legitimate if a package is installed with `yum install`. In Red Hat Enterprise Linux 7 the `--nogpgcheck` option is not required if a package is installed with `yum localinstall`.

```
[root@serverX ~]# yum localinstall
wonderwidgets-1.0-4.x86_64.rpm

[root@serverX ~]# rpm -q wonderwidgets
wonderwidgets-1.0-4.x86_64
```

### Note

`rpm -ivh PACKAGEFILE.rpm` can also be used to install package files. However, using `yum` helps maintain a transaction history kept by `yum` (see `yum history`).

### Warning

Be careful when installing packages from third parties, not just because of the software that they may install, but because the RPM may run arbitrary scripts as `root` as part of the installation process.

## Extracting files from RPM packages

Files in the RPM package can be extracted without installing the package using `cpio`, an archiving tool like `zip` or `tar`. Pipe the output of `rpm2cpio PACKAGEFILE.rpm` into `cpio -id`, and it will extract all the files stored in the RPM package. Subdirectory trees will be created as needed, relative to the current working directory.

Select files can also be extracted by specifying the path of the file:

```
[root@serverX ~]# rpm2cpio wonderwidgets-1.0-4.x86_64.rpm |
cpio -id "*txt"

11 blocks

[root@serverX ~]# ls -l usr/share/doc/wonderwidgets-1.0/
total 4
-rw-r--r--. 1 root root 76 Feb 13 19:27 README.txt
```

## Summary of rpm query commands

Installed packages can be queried directly with `rpm` command. Add a `-p` option to query a package file before installation.

| Task:                               | Command:                    |
|-------------------------------------|-----------------------------|
| Display information about a package | <code>rpm -q -i NAME</code> |

| Task:                                                        | Command:                             |
|--------------------------------------------------------------|--------------------------------------|
| List all files included in a package                         | <code>rpm -q -l NAME</code>          |
| List configuration files included in a package               | <code>rpm -q -c NAME</code>          |
| List documentation files included in a package               | <code>rpm -q -d NAME</code>          |
| Show a short summary of the reason for a new package release | <code>rpm -q --changelog NAME</code> |
| Display the shell scripts included in a package              | <code>rpm -q --scripts NAME</code>   |

#### References

`yum(8)`, `rpm(8)`, `repoquery(1)`, `rpm2cpio(8)`, and `cpio(1)` man pages

## 13.10. Practice: Working with RPM Package Files

### Guided exercise

In this lab, you will gather information about a third-party package, extract files from it, and install it on your serverX system.

#### Outcomes:

A package not provided by a yum repository is installed on the system.

#### Before you begin...

Reset your serverX system.

1. Download `wonderwidgets-1.0-4.x86_64.rpm` from <http://classroom/pub/materials>.

```

2. [root@serverX ~]# wget
 http://classroom/pub/materials/wonderwidgets-1.0-4.x
 86_64.rpm
3. --2014-02-11 15:58:02--
 http://classroom/pub/materials/wonderwidgets-1.0-4.x86_64.rpm
4. Resolving classroom (classroom)... 172.25.0.254
5. Connecting to classroom (classroom)|172.25.0.254|:80... connected.
6. HTTP request sent, awaiting response... 200 OK
7. Length: 5000 (4.9K) [application/x-rpm]
8. Saving to: 'wonderwidgets-1.0-4.x86_64.rpm'
9.
10. 100%[=====>] 5,000 --.-K/s in 0s
11.
12. 2014-02-11 15:58:02 (381 MB/s) - 'wonderwidgets-1.0-4.x86_64.rpm' saved
 [5000/5000]
```

13. What files does it contain?

```

14. [root@serverX ~]# rpm -q -p
 wonderwidgets-1.0-4.x86_64.rpm -l
```

```
15. /etc/wonderwidgets.conf
16. /usr/bin/wonderwidgets
17. /usr/share/doc/wonderwidgets-1.0
18. /usr/share/doc/wonderwidgets-1.0/README.txt
```

19. What scripts does it contain?

```
20. [root@serverX ~]# rpm -q -p
wonderwidgets-1.0-4.x86_64.rpm --scripts
```

21. How much disk space will it use when installed?

```
22. [root@serverX ~]# rpm -q -p
wonderwidgets-1.0-4.x86_64.rpm -i
23. Name : wonderwidgets
24. Version : 1.0
25. Release : 4
26. Architecture: x86_64
27. Install Date: (not installed)
28. Group : GLS/Applications
29. Size : 4849
30. License : GPL
31. Signature : (none)
32. Source RPM : wonderwidgets-1.0-4.src.rpm
33. Build Date : Fri 03 Dec 2010 05:42:55 AM EST
34. Build Host : station166.rosemont.lan
35. Relocations : (not relocatable)
36. Vendor : Red Hat, Inc.
37. Summary : Demonstration package for use in GLS training.
38. Description :
39. A demonstration package that installs an executable, and a config file.
```

40. Use **yum localinstall** to install the package.

```
41. [root@serverX ~]# yum localinstall
wonderwidgets-1.0-4.x86_64.rpm
42. Loaded plugins: langpacks
43. Examining wonderwidgets-1.0-4.x86_64.rpm: wonderwidgets-1.0-4.x86_64
44. Marking wonderwidgets-1.0-4.x86_64.rpm to be installed
45. Resolving Dependencies
```

```

46. --> Running transaction check
47. ---> Package wonderwidgets.x86_64 0:1.0-4 will be installed
48. --> Finished Dependency Resolution
49.
50. Dependencies Resolved
51.
52. =====
 ====
53. Package Arch Version Repository Size
54. =====
 ====
55. Installing:
56. wonderwidgets x86_64 1.0-4 /wonderwidgets-1.0-4.x86_64 4.7 k
57.
58. Transaction Summary
59. =====
 ====
60. Install 1 Package
61.
62. Total size: 4.7 k
63. Installed size: 4.7 k
64. Is this ok [y/d/N]: y
65. Downloading packages:
66. Running transaction check
67. Running transaction test
68. Transaction test succeeded
69. Running transaction
70. Installing : wonderwidgets-1.0-4.x86_64 1/1
71. Verifying : wonderwidgets-1.0-4.x86_64 1/1
72.
73. Installed:
74. wonderwidgets.x86_64 0:1.0-4
75.

```

Complete!

## 13.11. Lab: Installing and Updating Software Packages

### Performance checklist

In this lab, you will install and update select software packages.

#### Outcomes:

New and updated packages are installed on the system.

#### Before you begin...

Reset your serverX system.

1. Create the file `/etc/yum.repos.d/errata.repo`, to enable the “Updates” repository found on the content machine. It should access content found at the following URL:`http://content.example.com/rhel7.0/x86_64/errata`. Do *not* check GPG signatures.

Create the file `/etc/yum.repos.d/errata.repo` with the following content:

```
[updates]
name=Red Hat Updates
baseurl=http://content.example.com/rhel7.0/x86_64/errata
enabled=1
gpgcheck=0
```

2. Configure serverX to adhere to very specific software requirements. It must have the latest version of the following packages installed. Do not install all updates. Only install updates for the packages listed if they are available.
  - a. **kernel** (existing package w/ an update)  
`yum update kernel`
  - b. **xsane-gimp** (new package)  
`yum install xsane-gimp`
  - c. **rht-system** (new package)  
`yum install rht-system`
3. For security reasons, it should not have the **wvdial** package installed.  
`yum remove wvdial`
4. When you are ready to check your work, run **lab software grade** on serverX.

```
[student@serverX ~]$ lab software grade
```

# Chapter 14. Accessing Linux File Systems

14.1. Identifying File Systems and Devices

14.2. Practice: Identifying File Systems and Devices

14.3. Mounting and Unmounting File Systems

14.4. Practice: Mounting and Unmounting File Systems

14.5. Making Links Between Files

14.6. Practice: Making Links Between Files

14.7. Locating Files on the System

14.8. Practice: Locating Files on the System

14.9. Lab: Accessing Linux File Systems

## Introduction

| Overview          |                                                                                                                                                                                                                                                                               |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Goal</b>       | To access and inspect existing file systems on a Red Hat Enterprise Linux system.                                                                                                                                                                                             |
| <b>Objectives</b> | <ul style="list-style-type: none"><li>○ Identify the file system hierarchy.</li><li>○ Access the contents of file systems.</li><li>○ Use hard links and symlinks to make multiple names.</li><li>○ Search for files on mounted file systems.</li></ul>                        |
| <b>Sections</b>   | <ul style="list-style-type: none"><li>○ Identifying File Systems and Devices (and Practice)</li><li>○ Mounting and Unmounting File Systems (and Practice)</li><li>○ Making Links Between Files (and Practice)</li><li>○ Locating Files on the System (and Practice)</li></ul> |
| <b>Lab</b>        | <ul style="list-style-type: none"><li>○ Accessing Linux File Systems</li></ul>                                                                                                                                                                                                |



## 14.1. Identifying File Systems and Devices

### Objectives

After completing this section, students should be able to identify a directory in the file system hierarchy and what storage device it is stored on.

### Storage management concepts

A file system is an organized structure of data-holding files and directories residing on a storage device, such as a physical disk or partition. The file system hierarchy discussed earlier assembles all the file systems into one tree of directories with a single root, the `/` directory. The advantage here is that the existing hierarchy can be extended at any time by adding a new disk or partition containing a supported file system to add disk space anywhere in the file system tree. The process of adding a new file system to the existing directory tree is called *mounting*. The directory where the new file system is mounted is referred to as a *mount point*. This is a fundamentally different concept than that used on a Microsoft Windows system, where a new file system is represented by a separate drive letter.

Hard disks and storage devices are normally divided up into smaller chunks called *partitions*. A partition is a way to compartmentalize a disk. Different parts of it can be formatted with different file systems or used for different purposes. For example, one partition could contain user home directories while another could contain system data and logs. If a user fills up the home directory partition with data, the system partition may still have space available. Placing data in two separate file systems on two separate partitions helps in planning data storage.

Storage devices are represented by a special file type called *block device*. The block device is stored in the `/dev` directory. In Red Hat Enterprise Linux, the first **SCSI, PATA/SATA, or USB hard drive** detected is `/dev/sda`, the second is `/dev/sdb`, and so on. This name represents the whole drive. The first primary partition on `/dev/sda` is `/dev/sda1`, the second partition is `/dev/sda2`, and so on.

A long listing of the `/dev/vda` device file on serverX reveals its special file type as **b**, which stands for block device:

```
[student@serverX ~]$ ls -l /dev/vda
brw-rw----. 1 root disk 253, 0 Mar 13 08:00 /dev/vda
```

#### Note

Exceptions are hard drives in **virtual machines**, which typically show up as `/dev/vd<letter>` or `/dev/xvd<letter>`.

Another way of organizing disks and partitions is with *logical volume management (LVM)*. With LVM, one or more block devices **can be aggregated into a storage pool called a volume group**. Disk space is made available with one or **more logical volumes**. **A logical volume is the equivalent of a partition**

**residing** on a physical disk. Both the volume group and the logical volume have names assigned upon creation. For the volume group, a directory with the same name as the volume group exists in the `/dev` directory. Below that directory, a symbolic link with the same name as the logical volume has been created. For example, the device file representing the `mylv` logical volume in the `myvg` volume group is `/dev/myvg/mylv`.

It should be noted that LVM relies on the **Device Mapper (DM) kernel driver**. The above symbolic link `/dev/myvg/mylv` points to the `/dev/dm-number` block device node. The assignment of the *number* is **sequential beginning with zero (0)**. There is another symbolic link for every logical

属于

volume in the `/dev/mapper` directory with the name `/dev/mapper/myvg-mylv`. Access to the logical volume can generally use either of the consistent and reliable symbolic link names, since the `/dev/dm-number` name can vary with each boot.

## Examining file systems

### Examining file systems

To get an overview about the file system mount points and the amount of free space available, run the `df` command. When the `df` command is run without arguments, it will report total disk space, used disk space, and free disk space on all mounted regular file systems. It will report on both local and remote systems and the percentage of the total disk space that is being used. Display the file systems and mount points on the serverX machine.

```
[student@serverX ~]$ df
```

| Filesystem | 1K-blocks | Used    | Available | Use% | Mounted on     |
|------------|-----------|---------|-----------|------|----------------|
| /dev/vda1  | 6240256   | 4003760 | 2236496   | 65%  | /              |
| devtmpfs   | 950536    | 0       | 950536    | 0%   | /dev           |
| tmpfs      | 959268    | 80      | 959188    | 1%   | /dev/shm       |
| tmpfs      | 959268    | 2156    | 957112    | 1%   | /run           |
| tmpfs      | 959268    | 0       | 959268    | 0%   | /sys/fs/cgroup |

The partitioning on the serverX machine shows one real file system, which is mounted on `/`. This is common for virtual machines. The `tmpfs` and `devtmpfs` devices are file systems in system memory. All files written into `tmpfs` or `devtmpfs` disappear after system reboot.

To improve readability of the output sizes, there are two different *human-readable* options: `-h` or `-H`. The difference between these two options is that `-h` will report in KiB ( $2^{10}$ ), MiB ( $2^{20}$ ), or GiB ( $2^{30}$ ), while the `-H` option will report in SI units: KB ( $10^3$ ), MB ( $10^6$ ), GB ( $10^9$ ), etc. Hard drive manufacturers usually use SI units when advertising their products.

Show a report on the file systems on the serverX machine with all units converted to human-readable format:

```
[student@serverX ~]$ df -h
```

| Filesystem | Size | Used | Avail | Use% | Mounted on     |
|------------|------|------|-------|------|----------------|
| /dev/vda1  | 6.0G | 3.9G | 2.2G  | 65%  | /              |
| devtmpfs   | 929M | 0    | 929M  | 0%   | /dev           |
| tmpfs      | 937M | 80K  | 937M  | 1%   | /dev/shm       |
| tmpfs      | 937M | 2.2M | 935M  | 1%   | /run           |
| tmpfs      | 937M | 0    | 937M  | 0%   | /sys/fs/cgroup |

For more detailed information about the space used by a certain directory tree, there is the `du` command. The `du` command has `-h` and `-H` options to convert the output to human-readable format. The `du` command shows the size of all files in the current directory tree recursively.

Show a disk usage report for the `/root` directory on serverX:

```
[root@serverX ~]# du /root
```

4        /root/.ssh

```

4 /root/.cache/dconf
4 /root/.cache
4 /root/.dbus/session-bus
4 /root/.dbus
0 /root/.config/ibus/bus
0 /root/.config/ibus
0 /root/.config
14024 /root

```

Show a disk usage report in human-readable format for the `/var/log` directory on serverX:

```

[root@serverX ~]# du -h /var/log
...
4.9M /var/log/sa
68K /var/log/prelink
0 /var/log/qemu-ga
14M /var/log

```

## References

[df\(1\)](#) and [du\(1\)](#) man pages

## 14.2. Practice: Identifying File Systems and Devices

### Quiz

Match the following items to their counterparts in the table.

| Description                                                                                | Device file                          |
|--------------------------------------------------------------------------------------------|--------------------------------------|
| The device file of a SATA hard drive residing in <code>/dev.</code>                        | <code>/dev/sdc</code>                |
| The device file of the second partition on the first SATA hard drive in <code>/dev.</code> | <code>/dev/sda2</code>               |
| The device file of a logical volume in <code>/dev.</code>                                  | <code>/dev/vg_install/lv_home</code> |
| The device file of the second disk in a virtual machine in <code>/dev.</code>              | <code>/dev/vdb</code>                |
| The device file of the third partition on the second SATA hard drive in <code>/dev.</code> | <code>/dev/sdb3</code>               |
| The device file of the third partition on the second disk                                  | <code>/dev/vdb3</code>               |

| Description                                 | Device file |
|---------------------------------------------|-------------|
| in a virtual machine in <code>/dev</code> . |             |

## 14.3. Mounting and Unmounting File Systems

### Objectives

After completing this section, students should be able to access the contents of file systems by adding and removing file systems from the file system hierarchy.

### Mounting file systems manually

#### *Mounting and Unmounting File Systems*

A file system residing on a SATA/PATA or SCSI device needs to be mounted manually to access it.

The **mount** command allows the root user to manually mount a file system. The first argument of the **mount** command specifies the file system to mount. The second argument specifies the target directory where the file system is made available after mounting it. The target directory is referred to as a mount point.

The **mount** command expects the file system argument in one of two different ways:

- The device file of the partition holding the file system, residing in `/dev`.
- The *UUID*, a universal unique identifier of the file system.

#### Note

As long as a file system is not recreated, the UUID stays the same. The device file can change; for example, if the order of the devices is changed or if additional devices are added to the system.

The **blkid** command gives an overview of existing partitions with a file system on them and the UUID of the file system, as well as the file system used to format the partition.

```
[root@serverX ~]# blkid
/dev/vda1: UUID="46f543fd-78c9-4526-a857-244811be2d88" TYPE="xfs"
```

#### Note

A file system can be mounted on an existing directory. The `/mnt` directory exists by default and provides an entry point for mount points. It is used for manually mounting disks. It is recommended to create a subdirectory under `/mnt` and use that subdirectory as a mount point unless there is a reason to mount the file system in another specific location in the file system hierarchy.

Mount by device file of the partition that holds the file system.

```
[root@serverX ~]# mount /dev/vdb1 /mnt/mydata
```

Mount the file system by universal unique id, or the UUID, of the file system.

```
[root@serverX ~]# mount
UUID="46f543fd-78c9-4526-a857-244811be2d88" /mnt/mydata
```

#### Note

If the directory acting as mount point is not empty, the files that exist in that directory are not accessible as long as a file system is mounted there. All files written to the mount point directory end up on the file system mounted there.

## Unmounting file systems

To unmount a file system, the **umount** command expects the mount point as an argument. Change to the **/mnt/mydata** directory. Try to unmount the device mounted on the **/mnt/mydata** mount point. It will fail.

```
[root@serverX ~]# cd /mnt/mydata
[root@serverX mydata]# umount /mnt/mydata
umount: /mnt/mydata: target is busy.

(In some cases useful info about processes that use
the device is found by lsof(8) or fuser(1))
```

Unmounting is not possible if the mount point is accessed by a process. For **umount** to be successful, the process needs to stop accessing the mount point.

The **lsof** command lists all open files and the process accessing them in the provided directory. It is useful to identify which processes currently prevent the file system from successful unmounting.

```
[root@serverX mydata]# lsof /mnt/mydata
COMMAND PID USER FD TYPE DEVICE SIZE/OFF NODE NAME
bash 1593 root cwd DIR 253,2 6 128 /mnt/mydata
lsof 2532 root cwd DIR 253,2 19 128 /mnt/mydata
lsof 2533 root cwd DIR 253,2 19 128 /mnt/mydata
```

Once the processes are identified, an action can be taken, such as waiting for the process to complete or sending a SIGTERM or SIGKILL signal to the process. In this case, it is sufficient to change the current working directory to a directory outside the mount point.

```
[root@serverX mydata]# cd
[root@serverX ~]# umount /mnt/mydata
```

### Note

A common cause for the file system on the mount point to be busy is if the current working directory of a shell prompt is below the active mount point. The process accessing the mount point is **bash**. Changing to a directory outside the mount point allows the device to be unmounted.

## Accessing removable storage devices

Removable media, such as USB flash devices and drives, get automatically mounted by the graphical desktop environment when plugged in. The mount point for the removable medium is **/run/media/<user>/<label>**. The <user> is the user logged into the graphical environment. The <label> is the name given to the file system when it was created.

### Warning

To safely remove USB media from the system, it is required to unmount it before physically removing it from the USB slot to synchronize the file system. Removing a USB storage device without unmounting the file system on it can result in data loss.

## References

`mount(8)`, `umount(8)`, and `lsfs(8)` man pages

# 14.4. Practice: Mounting and Unmounting File Systems

## Guided exercise

In this lab, you will mount and unmount file systems.

### Outcomes:

The user identifies and mounts a new file system at a specified mount point, then unmounts it.

### Before you begin...

Reset your serverX system. Run the script `lab fs setup` before starting the exercise.

1. A new partition with a file system has been added to the second disk (vdb) on your serverX machine. Mount the newly available partition by UUID at the newly created mount point `/mnt/newspace`.

- a. Use the `blkid` to discover the UUID of the newly added partition, `vdb1`, on serverX.

```
b. [root@serverX ~]# blkid
c. /dev/xvda1: UUID="46f543fd-78c9-4526-a857-244811be2d88" TYPE="xfs"
```

```
/dev/xvdj1: UUID="7c5e3fbb-34eb-4431-a4a5-9b887c1b6866" TYPE="xfs"
```

- d. Create the mount point `/mnt/newspace` on serverX.

```
[root@serverX ~]# mkdir /mnt/newspace
```

- e. Mount the file system by UUID on the `/mnt/newspace` directory of the serverX machine.

```
[root@serverX ~]# mount
UUID="7c5e3fbb-34eb-4431-a4a5-9b887c1b6866"
/mnt/newspace
```

2. Change to the `/mnt/newspace` directory and create a new directory, `/mnt/newspace/newdir`, with an empty file, `/mnt/newspace/newdir/newfile`, on serverX.

- a. Change to the `/mnt/newspace` directory on serverX.

```
[root@serverX ~]# cd /mnt/newspace
```

- b. Create a new directory, `/mnt/newspace/newdir`, on serverX.

```
[root@serverX newspace]# mkdir newdir
```

- c. Create a new empty file, `/mnt/newspace/newdir/newfile`, on serverX.

```
[root@serverX newspace]# touch newdir/newfile
```

3. Unmount the file system mounted on the `/mnt/newspace` directory on serverX.
  - a. Try to unmount `/mnt/newspace` while the current directory on the shell is still `/mnt/newspace` on serverX.

```
[root@serverX newspace]# umount /mnt/newspace
```

- b. Change the current directory on the shell to `/root`.

```
c. [root@serverX newspace]# cd
```

```
[root@serverX ~]#
```

- d. Successfully unmount `/mnt/newspace` on serverX.

```
[root@serverX ~]# umount /mnt/newspace
```

## 14.5. Making Links Between Files

### Objectives

After completing this section, students should be able to use hard links and soft links to make multiple names point to the same file.

### Managing links between files

*Making links between files*

#### Creating hard links

A hard link is a new directory entry with a reference to an existing file on the file system. Every file in a file system has one hard link by default. To save space, instead of copying, a new hard link can be created to reference the same file. A new hard link either needs to have a different file name, if it is created in the same directory as the existing hard link, or it needs to reside in a different directory. All hard links pointing to the same file have the same permissions, link count, user/group ownerships, time stamps, and file content. Hard links pointing to the same file content need to be on the same file system.

The `ls -l` shows the hard link count after the permissions and before the owner of a file.

```
[root@serverX ~]# echo "Hello World" > newfile.txt
[root@serverX ~]# ls -l newfile.txt
-rw-r--r--. 1 root root 0 Mar 11 19:19 newfile.txt
```

The command `ln` creates new hard links to existing files. The command expects an existing file as the first argument, followed by one or more additional hard links. The hard links can reside anywhere as

long as they are on the same file system as the existing file. After a new hard link is created, there is no way to tell which of the existing hard links is the original one.

Create a hard link `newfile-link2.txt` for the existing file `newfile.txt` in the same directory.

```
[root@serverX ~]# ln newfile.txt /tmp/newfile-hlink2.txt
[root@serverX ~]# ls -l newfile.txt /tmp/newfile-hlink2.txt
-rw-rw-r--. 2 root root 12 Mar 11 19:19 newfile.txt
-rw-rw-r--. 2 root root 12 Mar 11 19:19 newfile-hlink2.txt
```

Even if the original file gets deleted, the content of the file is still available as long as at least one hard link exists.

```
[root@serverX ~]# rm -f newfile.txt
[root@serverX ~]# ls -l /tmp/newfile-link2.txt
-rw-rw-r--. 1 root root 12 Mar 11 19:19 /tmp/newfile-link2.txt
[root@serverX ~]# cat /tmp/newfile-link2.txt
Hello World
```

### Important

All hard links referencing the same file have the same permissions, link count, user/group ownerships, time stamps, and file content. If any of that information is changed on one hard link, all other hard links pointing at the same file will show the new information as well.

### Creating soft links

The `ln -s` command creates a soft link, also known as a symbolic link. A soft link is a special file type that points to an existing file or directory. Soft links can point to a file or directory on another file system. Unlike a hard link, a symbolic link can point to a file on a different file system.

```
[root@serverX ~]# ln -s /root/newfile-link2.txt
/tmp/newfile-symlink.txt
[root@serverX ~]# ls -l newfile-link2.txt
/tmp/newfile-symlink.txt
lrwxrwxrwx. 1 root root 11 Mar 11 20:59 /tmp/newfile-symlink.txt ->
/root/newfile-link2.txt
-rw-rw-r--. 1 root root 12 Mar 11 19:19 newfile-link2.txt
```

When the original file gets deleted, the soft link is still pointing to the file but the target is gone. A soft link pointing to a missing file is called a "dangling soft link."

```
[root@serverX ~]# rm -f newfile-link2.txt
[root@serverX ~]# ls -l /tmp/newfile-symlink.txt
lrwxrwxrwx. 1 root root 11 Mar 11 20:59 /tmp/newfile-symlink.txt ->
newfile-link2.txt
[root@serverX ~]# cat /tmp/newfile-symlink.txt
```



```
cat: /tmp/newfile-symlink.txt: No such file or directory
```

A soft link can point to a directory. The soft link then acts like a directory. Changing to the soft link directory with **cd** works as expected.

Create a soft link **/root/configfiles** pointing to the **/etc** directory.

```
[root@serverX ~]# ln -s /etc /root/configfiles
[root@serverX ~]# cd /root/configfiles
[root@serverX configfiles]# pwd
/root/configfiles
```

## References

[ln\(1\) man page](#)

## 14.6. Practice: Making Links Between Files

### Guided exercise

In this lab, you will create hard and soft links.

#### Outcomes:

The user creates a hard link and a soft link.

1. Create an additional hard link **/root/qmp-manual.txt** for the existing file **/usr/share/doc/qemu-kvm/qmp-commands.txt** on serverX.
  - a. Create the hard link **/root/qmp-manual.txt**. Link it to the file **/usr/share/doc/qemu-kvm/qmp-commands.txt**.

```
[root@serverX ~]# ln
/usr/share/doc/qemu-kvm/qmp-commands.txt
/root/qmp-manual.txt
```

- b. Verify the link count on the newly created link **/root/qmp-manual.txt**.

```
c. [root@serverX ~]# ls -l /root/qmp-manual.txt
```

```
-rw-r--r--. 2 root root 63889 Nov 11 02:58 /root/qmp-manual.txt
```

- d. Verify the link count on the original file **/usr/share/doc/qemu-kvm/qmp-commands.txt**.

```
e. [root@serverX ~]# ls -l
/usr/share/doc/qemu-kvm/qmp-commands.txt
```

```
-rw-r--r--. 2 root root 63889 Nov 11 02:58
/usr/share/doc/qemu-kvm/qmp-commands.txt
```

2. Create the soft link **/root/tempdir** pointing to the directory **/tmp** on serverX.
  - a. Create the soft link **/root/tempdir**. Link it to **/tmp**.

```
[root@serverX ~]# ln -s /tmp /root/temdir
```

b. Verify the newly created link with `ls -l`.

```
c. [root@serverX ~]# ls -l /root
```

```
lrwxrwxrwx. 1 root root 4 Mar 13 08:42 /root/temdir -> /tmp
```

## 14.7. Locating Files on the System

### Objectives

After completing this section, students should be able to search for files on mounted file systems using `find` and `locate`.

### Tools for finding files

A system administrator needs tools for searching files matching certain criteria on the file system. This section discusses two commands that can search files in the file system. The `locate` command searches a pregenerated database for file names or file paths and returns the results instantly. The `find` command searches the file system in real time by crawling through the file system.

### Locating files by name with `locate`

*Locating files by name with `locate`*

The `locate` command returns search results based on file name or path from the `locate` database. The database stores file name and path information.

When searching entries as a regular user, results are returned only for where the user invoking the `locate` search has read permissions on the directory tree that contains the matching element. Search for files with "passwd" in the name or path in directory trees readable by user student on serverX.

```
[student@serverX ~]$ locate passwd
```

```
/etc/passwd
```

```
/etc/passwd-
```

```
/etc/pam.d/passwd
```

```
/etc/security/opasswd
```

```
/usr/bin/gpasswd
```

```
/usr/bin/grub2-mkpasswd-pbkdf2
```

```
/usr/bin/lpasswd
```

```
/usr/bin/passwd
```

```
/usr/bin/userpasswd
```

```
/usr/bin/vino-passwd
```

```
/usr/bin/vncpasswd
```

Results are returned even when the file name or path is only a partial match to the search query.

```
[root@serverX ~]# locate image
/home/myuser/boot.image
/home/someuser/my_family_image.png
/home/student/myimages-vacation/picture.png
```

The **-i** option performs a case-insensitive search. With this option, all possible combinations of upper- and lowercase letters match the search.

```
[student@serverX ~]$ locate -i messages
...
/usr/share/vim/vim74/lang/zh_TW/LC_MESSAGES
/usr/share/vim/vim74/lang/zh_TW/LC_MESSAGES/vim.mo
/usr/share/vim/vim74/lang/zh_TW.UTF-8/LC_MESSAGES
/usr/share/vim/vim74/lang/zh_TW.UTF-8/LC_MESSAGES/vim.mo
/usr/share/vim/vim74/syntax/messages.vim
/usr/share/vim/vim74/syntax/msmessages.vim
/var/log/messages
```

The **-n** option limits the number of returned search results by **locate**. The following example limits the search results returned by **locate** to the first five matches.

```
[student@serverX ~]$ locate -n 5 snow.png
/usr/share/icons/HighContrast/16x16/status/weather-snow.png
/usr/share/icons/HighContrast/22x22/status/weather-snow.png
/usr/share/icons/HighContrast/24x24/status/weather-snow.png
/usr/share/icons/HighContrast/256x256/status/weather-snow.png
/usr/share/icons/HighContrast/32x32/status/weather-snow.png
```

#### Note

The **locate** database is automatically updated every day. The root user can perform an update of the database with the **updatedb** command.

```
[root@serverX ~]# updatedb
```

## ▣ Searching for files with find

### Searching for files with *find*

The **find** command performs a real-time search in the local file systems to find files that match the criteria of the command-line arguments. The **find** command is looking at files in the file system as your user account. The user invoking the **find** command must have read and execute permission on a directory to examine its contents.

The first argument to the **find** command is the directory to search. If the directory argument is omitted, **find** will start the search in the current directory and look for matches in any of the subdirectories.

To search the home directory of user student, give **find** a starting directory of **/home/student**. To search the entire system, provide a starting directory of **/**.

#### Note

**find** has a huge number of options that can be provided to describe exactly what kind of file should be found. Searches can be based on file name, file size, last modified time stamp, and other file characteristics in any combination.

The **-name** option followed by a file name looks up files matching the given file name and returns all exact matches. To search for files named **sshd\_config** in the **/** directory and all subdirectories on serverX, run:

```
[root@serverX ~]# find / -name sshd_config
/etc/ssh/sshd_config
```

Wild cards are available to search for a file name and return all results that are a partial match. When using wild cards, it is important to quote the file name to look for to prevent the terminal from interpreting the wild card.

The following example searches for files in the **/** directory on serverX that end in **.txt**:

```
[root@serverX ~]# find / -name '*.txt'
/etc/pki/nssdb/pkcs11.txt
/etc/brltty/brl-lt-all.txt
/etc/brltty/brl-mb-all.txt
/etc/brltty/brl-md-all.txt
/etc/brltty/brl-mn-all.txt
...
```

To search for files in **/etc/** that contain **pass** anywhere in their names on serverX, run:

```
[root@serverX ~]# find /etc -name '*pass*'
/etc/passwd
/etc/passwd-
/etc/fonts/conf.d/60-overpass.conf
/etc/selinux/targeted/modules/active/modules/passenger.pp
/etc/security/opasswd
/etc/pam.d/passwd
/etc/pam.d/password-auth-ac
/etc/pam.d/password-auth
/etc/pam.d/gdm-password
```

To perform a case-insensitive search for a given file name, use the **-iname** option, followed by the file name to search. To search case-insensitively for files that have **messages** in their names in the **/** directory on serverX, run:

```
[root@serverX ~]# find / -iname '*messages*'
/var/log/messages
/usr/lib64/python2.7/site-packages/orca/notification_messages.py
/usr/lib64/python2.7/site-packages/orca/notification_messages.pyc
/usr/lib64/python2.7/site-packages/orca/notification_messages.pyo
/usr/share/locale/aa/LC_MESSAGES
/usr/share/locale/ab/LC_MESSAGES
/usr/share/locale/ace/LC_MESSAGES
```

**find** can search for files based on their ownership or permissions. Useful options when searching by owner are **-user** and **-group**, which search by name, and **-uid** and **-gid**, which search by ID.

Search for files owned by the user *student* in the **/home/student** directory on serverX.

```
[student@serverX ~]$ find -user student
.
./.bash_logout
./.bash_profile
./.bashrc
./.ssh
...
```

Search for files owned by the group *student* in **/home/** in the **/home/student** directory on serverX.

```
[student@serverX ~]$ find -group student
.
./.bash_logout
./.bash_profile
./.bashrc
./.ssh
...
```

Search for files owned by user ID *1000* in the **/home/student** directory on serverX.

```
[student@serverX ~]$ find -uid 1000
.
./.bash_logout
./.bash_profile
```

```
./bashrc
./ssh
...
```

Search for files owned by group ID *1000* in the **/home/student** directory on serverX.

```
[student@serverX ~]$ find -gid 1000
.
./bash_logout
./bash_profile
./bashrc
./ssh
...
```

Search for files owned by user *root* and group *mail* on the serverX machine.

```
[root@serverX ~]# find / -user root -group mail
/var/spool/mail
/var/spool/mail/root
```

The **-perm** option is used to look for files with a particular set of permissions. Permissions can be described as octal values, with some combination of 4, 2, and 1 for read, write, and execute.

Permissions can be preceded by a **/** or **-** sign.

A numeric permission preceded by **/** will match files that have at least one bit of user, group, or other for that permission set. A file with permissions **r--r--r--** does not match **/222**, but one with **rw-r--r--** does. A **-** sign before a permission means that all three instances of that bit must be on, so neither of the previous examples would match, but something like **rw-rw-rw-** would.

To use a more complex example, the following command would match any file for which the user has read, write, and execute permissions, members of the group have read and write permissions, and others have read-only access:

```
[root@serverX ~]# find /home -perm 764
```

To match files for which the user has at least write and execute permissions, *and* the group has at least write permissions, *and* others have at least read access:

```
[root@serverX ~]# find /home -perm -324
```

To match files for which the user has read permissions, *or* the group has at least read permissions, *or* others have at least write access:

```
[root@serverX ~]# find /home -perm /442
```

When used with **/** or **-**, a value of **0** works like a wild card, since it means "a permission of at least nothing."

To match any file in the **/home/student** directory for which others have at least read access on serverX, run:

```
[student@serverX ~]$ find -perm -004
```

Find all files in the `/home/student` directory where *other* has write permissions on serverX.

```
[student@serverX ~]$ find -perm -002
```

The **find** command can look up files that match a size specified with the **-size** option, followed by a numerical value and the unit.

Units to be used with the **-size** option are:

- k, for kilobyte
- M, for megabyte
- G, for gigabyte

Search for files with a size of *exactly* 10 megabytes.

```
[student@serverX ~]$ find -size 10M
```

Find files with a size *more than* 10 gigabytes.

```
[student@serverX ~]$ find -size +10G
```

List all files with a size *less than* 10 kilobytes.

```
[student@serverX ~]$ find -size -10k
```

#### Important

The **-size** unit modifiers round everything up to single units. For example, **find -size 1M** will show files smaller than 1MB because it rounds all files up to 1MB.

The **-mmin** option, followed by the time in minutes, searches for all files that had their content changed at exactly the given time in the past.

To find all files that had their file content changed exactly 120 minutes ago on serverX, run:

```
[root@serverX ~]# find / -mmin 120
```

The **+** modifier in front of the amount of minutes looks for all files in the `/` that have been modified more than 200 minutes ago.

```
[root@serverX ~]# find / -mmin +200
```

The **-** modifier changes the search to look for all files in the `/` directory which have been changed less than 150 minutes ago.

```
[root@serverX ~]# find / -mmin -150
```

The **-type** option limits the search scope to a given file type, such as:

- **f**, for regular file
- **d**, for directory
- **l**, for soft link
- **b**, for block device

Search for all directories in the `/etc` folder on serverX.

```
[root@serverX ~]# find /etc -type d
/etc
/etc/tmpfiles.d
/etc/systemd
/etc/systemd/system
/etc/systemd/system/getty.target.wants
...
```

Search for all soft links on the serverX system.

```
[root@serverX ~]# find / -type l
```

Generate a list of all block devices in the `/dev` directory on serverX:

```
[root@serverX ~]# find /dev -type b
/dev/vda1
/dev/vda
```

The `-links` option followed by a number looks for all files that have a certain hard link count. The number can be preceded by a `+` modifier to look for files with a count higher than the given hard link count. If the number is preceded with a `-` modifier, the search is limited to all files with a hard link count that is less than the given number.

Search for all regular files with more than one hard link on the serverX machine:

```
[root@serverX ~]# find / -type f -links +1
```

## References

`locate`(1), `updatedb`(8), and `find`(1) man pages

## 14.8. Practice: Locating Files on the System

### Guided exercise

In this lab, students will find files on the local file system.

#### Outcomes:

The user will search files with `locate` and `find`.

1. Use the `locate` command to find various different files on the serverX machine.
  - a. Even though the `locate` database is updated automatically every day, make sure the database is up-to-date by manually starting an update on serverX.

```
[root@serverX ~]# updatedb
```

- b. Locate the configuration file `logrotate.conf` on serverX.

- c. 

```
[root@serverX ~]# locate logrotate.conf
```
- d. 

```
/etc/logrotate.conf
```



```
/usr/share/man/man5/logrotate.conf.5.gz
```

- e. Locate the configuration file **networkmanager.conf**, ignoring case, on serverX.

```
f. [root@serverX ~]# locate -i networkmanager.conf
g. /etc/NetworkManager/NetworkManager.conf
h. /etc/dbus-1/system.d/org.freedesktop.NetworkManager.conf
```

```
/usr/share/man/man5/NetworkManager.conf.5.gz
```

2. Use the **find** command to perform real-time searches on the serverX machine according to the following requirements:

- a. Find all files in the **/var/lib** directory owned by user **chrony** on serverX.

```
b. [root@serverX ~]# find /var/lib -user chrony
```

```
/var/lib/chrony
```

- c. List all files in the **/var** directory owned by user **root** and group **mail**.

```
d. [root@serverX ~]# find /var -user root -group mail
```

```
e. /var/spool/mail
```

```
/var/spool/mail/root
```

- f. List all files in the **/usr/bin** directory with a file size greater than 50 kilobytes.

```
g. [root@serverX ~]# find /usr/bin -size +50k
```

```
h. /usr/bin/pre-grohtml
```

```
i. /usr/bin/iconv
```

```
j. /usr/bin/localedef
```

```
k. /usr/bin/rpcgen
```

```
l. /usr/bin/less
```

```
...
```

- m. Find all files in the **/home/student** directory that have not been changed in the last 120 minutes on serverX.

```
n. [root@serverX ~]# find /home/student -mmin +120
```

```
o. /home/student
```

```
p. /home/student/.bash_logout
```

```
q. /home/student/.bash_profile
r. /home/student/.bashrc
s. /home/student/.ssh
```

...

- t. Find all files in the `/tmp` directory that have been changed in the last 240 minutes on serverX.

```
[root@serverX ~]# find /tmp -mmin -240
v. /tmp
w. /tmp/.X11-unix
x. /tmp/.X11-unix/X0
y. /tmp/.ICE-unix
```

...

## 14.9. Lab: Accessing Linux File Systems

### Performance checklist

In this lab, students will mount a local file system, review it, and work with soft links.

#### Outcomes:

- Generate a disk usage report.
- Mount a file system.
- Create a soft link.
- Search files in the local file system.

#### Before you begin...

Reset your serverX system.

Run the **lab fs setup** to set up the server machine for your exercise.

1. Generate a disk usage report with the **du** command of the `/var/log` directory on serverX, and save the result in the `/tmp/results.txt` file.

```
[root@serverX ~]# du /var/log >/tmp/results.txt
```

2. Identify and mount a newly added file system by UUID on the `/mnt/myfreespace` directory on serverX.
  - a. Identify the newly added file system with the **blkid** command on serverX.
  - b. Create the mount point `/mnt/myfreespace` on serverX.

```
[root@serverX ~]# mkdir /mnt/myfreespace
```

c. Mount the file system by UUID on the `/mnt/myfreespace` directory of the serverX machine.

```
[root@serverX ~]# mount
UUID="a84f6842-ec1d-4f6d-b767-b9570f9fcdc0"
/mnt/myfreespace
```

3. Create the soft link `/root/myfreespace`, which points to the `/mnt/myfreespace` directory on serverX.

```
[root@serverX ~]# ln -s /mnt/myfreespace
/root/myfreespace
```

4. Find all soft links on serverX that have `freespace` as part of their names.

```
5. [root@serverX ~]# find / -type l -name '*freespace*'
```

```
/root/myfreespace
```

# Chapter 15. Using Virtualized Systems

15.1. Managing a Local Virtualization Host

15.2. Practice: Managing a Local Virtualization Host

15.3. Installing a New Virtual Machine

15.4. Practice: Installing a New Virtual Machine

15.5. Chapter Test: Using Virtualized Systems

Introduction

| Overview          |                                                                                                                                                                                                                                |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Goal</b>       | To create and use Red Hat Enterprise Linux virtual machines with <b>Kernel-based Virtual Machine (KVM)</b> and libvirt.                                                                                                        |
| <b>Objectives</b> | <ul style="list-style-type: none"><li>○ Install a Red Hat Enterprise Linux system as a host for running virtual machines.</li><li>○ Perform an interactive install of Red Hat Enterprise Linux on a virtual machine.</li></ul> |
| <b>Sections</b>   | <ul style="list-style-type: none"><li>○ Managing a Local Virtualization Host (and Practice)</li><li>○ Installing a New Virtual Machine (and Practice)</li></ul>                                                                |

## Overview

|         |                             |
|---------|-----------------------------|
| Chapter | o Using Virtualized Systems |
| Test    |                             |

# 15.1. Managing a Local Virtualization Host

## Objectives

After completing this section, students should be able to:

- o Describe and **对比** **contrast** Red Hat virtualization platforms.
- o Install Red Hat Enterprise Linux as a virtualization host system.

## System virtualization and Red Hat Enterprise Linux

Managing a local virtualization host

**KVM (Kernel-based Virtual Machine) is a full virtualization solution** built into the standard Red Hat Enterprise Linux kernel. It can run multiple, unmodified Windows and Linux guest operating systems. The KVM hypervisor in Red Hat Enterprise Linux is managed with the **libvirt API** and utilities, such as **virt-manager** and **virsh**. Since Red Hat Enterprise Linux is the foundation of Red Hat Enterprise Virtualization and the **Red Hat OpenStack platform**, KVM is a consistent component across products of the Red Hat **cloud infrastructure**. **云架构的支持**

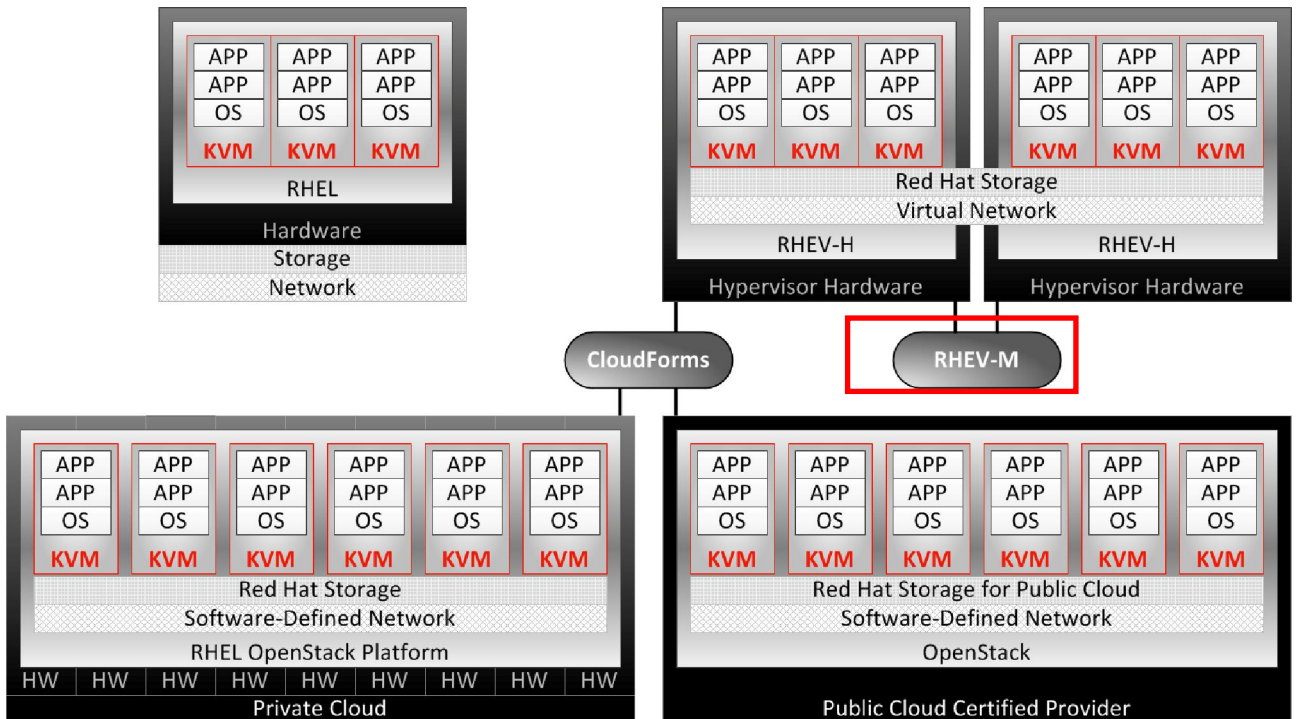


Figure 15.1: KVM throughout the Red Hat cloud infrastructure

KVM provides the virtual machine (VM) technology across all Red Hat products, from standalone physical instances of Red Hat Enterprise Linux up to the cloud OpenStack platform. Starting at the

## 虚拟化管理工具

upper-left corner of the previous figure:

- **Physical (legacy) systems**—Red Hat Enterprise Linux installations on legacy hardware provide KVM virtualization, up to the physical limitations of single systems, and are managed by libvirt utilities such as **virt-manager**. Red Hat Enterprise Linux instances may also be directly hosted at Red Hat Certified Cloud Providers through Red Hat Cloud Access.

Red Hat Enterprise Linux is typically configured as a thick host, a system that **supports VMs** while also providing other local and network services, applications, and management functions.

- **Red Hat Enterprise Virtualization (RHEV)**—supports KVM instances across multiple Red Hat Enterprise Virtualization hypervisor (**RHEV-H systems**), providing KVM **migration**, redundancy, and high availability managed by RHEV Manager (RHEV-M).

RHEV-H提供了KVM的迁移、冗余、和高可用性管理

Red Hat Enterprise Virtualization Hypervisor is a thin host, an expertly minimized and tuned version of Red Hat Enterprise Linux dedicated to the singular purpose of provisioning and supporting guest VMs.

- **RHEL OpenStack platform**—Red Hat private cloud architecture using integrated and tuned OpenStack on a Red Hat Enterprise Linux foundation with KVM, managed by either the Red Hat OpenStack Dashboard (Horizon component) or by Red Hat CloudForms.
- **OpenStack in public cloud**—OpenStack public cloud architecture implemented at Red Hat Certified Cloud Providers, and managed by either the OpenStack Horizon component or by Red Hat CloudForms.
- **Hybrid Cloud**—Red Hat CloudForms cloud management utilities manage and migrate KVM instances across Red Hat RHEV and OpenStack architectures, and transition KVM instances with third-party OpenStack and VMware platforms.

KVM instance configurations are compatible across Red Hat products. Installation requirements, parameters, and procedures are equivalent on supported platforms.

### Configure a Red Hat Enterprise Linux Physical System as a Virtualization Host

恰当的部署  
可以用于测试、实验

Red Hat Enterprise Linux can be configured as a virtualization host, **appropriate for development, testing, training**, or when needing to work in multiple operating systems at the same time. Red Hat Enterprise Linux hosts provide the ability to install additional software on the host platform as needed, such as **monitoring utilities** and **agents**, **network services**, **specialized storage**, and/or **other development tools** that may not be appropriate to install on dedicated Red Hat Enterprise Virtualization hypervisors.

调节和资源管理工具

Red Hat Enterprise Linux installations also provide easier access to **tuning and resource management** tools (e.g., **tuned and cgroups**). By comparison, **RHEV-H hypervisors** are highly secured and self-tuned, restricting system administrator-initiated customization by design. When greater administrative control is required and the performance compromise is acceptable, Red Hat Enterprise Linux is a **flexible standalone KVM platform**. KVM instances built on RHEL can be migrated or transitioned to more appropriate KVM platforms as enterprise needs increase.

灵活的管理平台，  
可以在多个KVM上  
迁移和传输

Preparing a Red Hat Enterprise Linux system to become a virtualization host requires checking for minimal system requirements and installing a choice of virtualization host packages.

Recommended system requirements:

- One processor core or hyperthread for the maximum number of **virtualized** CPUs in a guest virtual machine and one for the host.
- **2GB of RAM**, plus additional RAM for virtual machines.
- **6GB disk** space for the host, plus the required disk space for each virtual machine. Most guest operating systems require at least 6GB of disk space, but actual storage space requirements depend on each guest's image format.

The KVM hypervisor 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 that the host system hardware supports the correct extensions, view `/proc/cpuinfo`.

```
[root@serverX ~]# grep --color -E "vmx|svm" /proc/cpuinfo
```

```
flags : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat
pse36 clflush dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx lm constant_tsc
arch_perfmon pebs bts rep_good aperfmperf pni dtes64 monitor ds_cpl vmx smx est
tm2 ssse3 cx16 xtpr pdcm sse4_1 xsave lahf_lm dts tpr_shadow vnmi flexpriority
```

The No eXecute (NX) feature, called eXecute Disable (XD) by Intel and Enhanced Virus Protection by AMD, is not necessary for building a host on Red Hat Enterprise Linux, but is required for a Red Hat Enterprise Virtualization hypervisor (RHEV-H).

```
[root@serverX ~]# grep --color -E "nx" /proc/cpuinfo
```

```
flags : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat
pse36 clflush dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx lm constant_tsc
arch_perfmon pebs bts rep_good aperfmperf pni dtes64 monitor ds_cpl vmx smx est
tm2 ssse3 cx16 xtpr pdcm sse4_1 xsave lahf_lm dts tpr_shadow vnmi flexpriority
```

Building a RHEL virtualization host requires the **qemu-kvm** and **qemu-img** packages at minimum, to provide the **user-level KVM emulator** and **disk image manager**.

```
[root@serverX ~]# yum install qemu-kvm qemu-img
```

**Additional virtualization management packages** are also recommended:

- **python-virtinst**—Provides the virt-install command for creating virtual machines.

- **libvirt**—Provides the **host and server libraries** for interacting with hypervisors and host systems.
- **libvirt-python**—Contains a module to permit Python applications to use the libvirt API.
- **virt-manager**—Provides the Virtual Machine Manager graphical tool for administering VMs, using the libvirt-client library as the management API.
- **libvirt-client**—Provides the **client APIs and libraries for accessing libvirt servers**, including the virsh command-line tool to manage and control VMs.

```
[root@serverX ~]# yum install virt-manager libvirt
libvirt-python python-virtinst libvirt-client
```

The updated **anaconda** graphical installation program for Red Hat Enterprise Linux 7 provides better support for installing RHEL to meet specific purposes. An **anaconda** install no longer provides the ability to select individual RPM packages—only base environments and add-ons appropriate for the selected base—thus eliminating guesswork and resulting in leaner configurations. System administrators may still install any other desired RPM packages after an installation is complete by using standard RPM installation tools (e.g., **yum** or GNOME PackageKit).

To build a virtualization host during a graphical install of Red Hat Enterprise Linux, choose the base environment **Virtualization Host** in the left pane of the **anaconda Software Selection** screen. Select the **Virtualization Platform Add-On** checkbox in the right pane to include the management tools and utilities, as shown in the following figure.

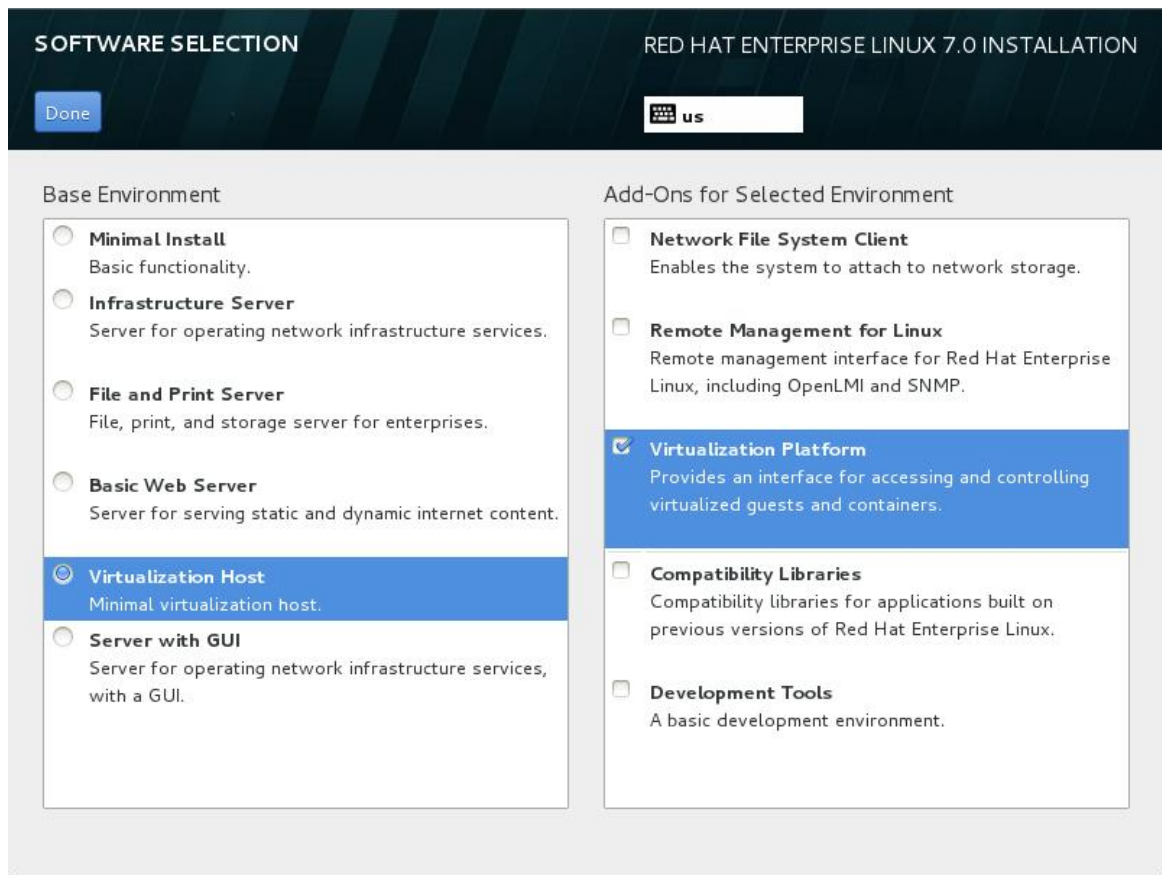


Figure 15.2: Creating a virtualization host during a graphical installation

## Managing virtual machines

The libvirt package is a hypervisor-independent virtualization API to securely manage virtual machines by providing the ability to provision, create, modify, monitor, control, migrate, and stop virtual machines on a single host. The libvirt package provides APIs to enumerate, monitor, and use the resources available on the managed host, including CPUs, memory, storage, and networking. Management tools utilizing libvirt can access host systems remotely using secure protocols.

Red Hat Enterprise Linux uses libvirt-based tools as the default for virtualization management. Support is included for the RHEL 5 Xen hypervisor, and for KVM on RHEL 5, 6, and 7. These management tools utilize libvirt:

- **virsh**—The virsh command-line tool is an alternative to the graphical virt-manager application. Unprivileged users can use virsh in read-only mode, or with root access for full administrative functionality. The virsh command is ideal for scripting virtualization administration.
- **virt-manager**—virt-manager is a graphical desktop tool allowing access to guest consoles, and is used to perform virtual machine creation, migration, configuration, and administrative tasks. Both local and remote hypervisors can be managed through a single interface.
- **RHEV-M**—Red Hat Enterprise Virtualization Manager provides a central management platform for physical and virtual resources, allowing virtual machines to be started, stopped, built, and migrated between hosts. RHEV-M also manages the storage and network components of a data center and provides secure, remote graphical guest console access.

Start the Virtual Machine Manager from the menu at **Applications** → **System Tools** → **Virtual Machine**



**Manager**, or by running the **virt-manager** command from the shell. Use this interface to power on or power off virtual machines, assign memory and CPU resources, monitor performance, and connect to the console of the virtual machines.

The **virsh** command-line tool provides the same functionality as **virt-manager**. Use **virsh** as an interactive shell to perform subcommands such as edit, list, start, stop, and destroy. The following examples show the **virsh** commands run as standalone commands from the shell:

```
[root@foundationX ~]# virsh list
```

| Id    | Name    | State   |
|-------|---------|---------|
| ----- |         |         |
| 1     | desktop | running |
| 2     | server  | running |

```
[root@foundationX ~]# virsh destroy server
```

```
[root@foundationX ~]# virsh list --all
```

| Id    | Name    | State    |
|-------|---------|----------|
| ----- |         |          |
| 1     | desktop | running  |
| -     | server  | shut off |

```
[root@foundationX ~]# virsh start server
```

```
[root@foundationX ~]# virsh list
```

| Id    | Name    | State   |
|-------|---------|---------|
| ----- |         |         |
| 1     | desktop | running |
| 2     | server  | running |

**virsh** has subcommands for additional management tasks:

- **connect**—Connect to a local or remote KVM host using **qemu:///host** syntax.
- **nodeinfo**—Returns basic information about the host, including CPUs and memory.
- **autostart**—Configures a KVM domain to start when the host boots.
- **console**—Connect to the virtual serial console of a guest.
- **create**—Create a domain from an XML configuration file and start it.
- **define**—Create a domain from an XML configuration file, but do not start it.
- **undefine**—Undefine a domain. If the domain is inactive, the domain configuration is removed.
- **edit**—Edit the XML configuration file for a domain, which will affect the next boot of the guest.
- **reboot**—Reboot the domain, as if the **reboot** command had been run from inside the guest.
- **shutdown**—Gracefully shuts down the domain, as if the **shutdown** command had been run from inside the guest.
- **screenshot**—Takes a screenshot of a current domain console and stores it in a file.

#### References

Additional information may be available in the introduction and chapter on system requirements in the Red Hat Enterprise Linux Virtualization Deployment and Administration Guide for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

Red Hat Enterprise Virtualization Administration Guide

- Section 1. Basics

Red Hat Enterprise Linux OpenStack Platform 4 Getting Started Guide

- Section 1. Introduction

**virsh(1)**, **virt-manager(1)** man pages

## ▮ 15.2. Practice: Managing a Local Virtualization Host

### Quiz

Match the following items to their counterparts in the table.

| Purpose                                                       | virsh subcommand |
|---------------------------------------------------------------|------------------|
| Boot an existing configured virtual machine                   | start            |
| Immediately stop a virtual machine, similar to unplugging it  | destroy          |
| Delete the configuration for a virtual machine permanently    | undefine         |
| Use an XML configuration to create and boot a virtual machine | create           |
| Use an XML configuration to create a virtual machine          | define           |
| Gracefully stop and restart a virtual machine                 | reboot           |
| Gracefully stop a virtual machine                             | shutdown         |

## 15.3. Installing a New Virtual Machine

### Objectives

After completing this section, students should be able to:

- Build a virtual machine configuration.
- Install Red Hat Enterprise Linux into a new virtual machine instance.

### Creating a virtual machine

#### *Installing a new virtual machine*

Virtual machines can perform the same tasks as physical systems. System administrators make system sizing and configuration decisions using fundamentally the same criteria as for physical systems, including the guest machine's role and projected system load. Preparatory considerations must include CPU and memory requirements, type of I/O and the expected number of clients, access to public or exclusive storage, current and future sizing expectations, and bandwidth and latency requirements. Necessary disk and network components must be configured on the virtual host before creating the guest.

Launch the Virtual Machine Manager from the menu **Applications** → **System Tools** → **Virtual Machine Manager**, or run the **virt-manager** command as root. Click the **Create a new virtual machine** button to open the **New VM** wizard. Ensure that virt-manager can access the installation media (whether locally or over the network) before continuing.

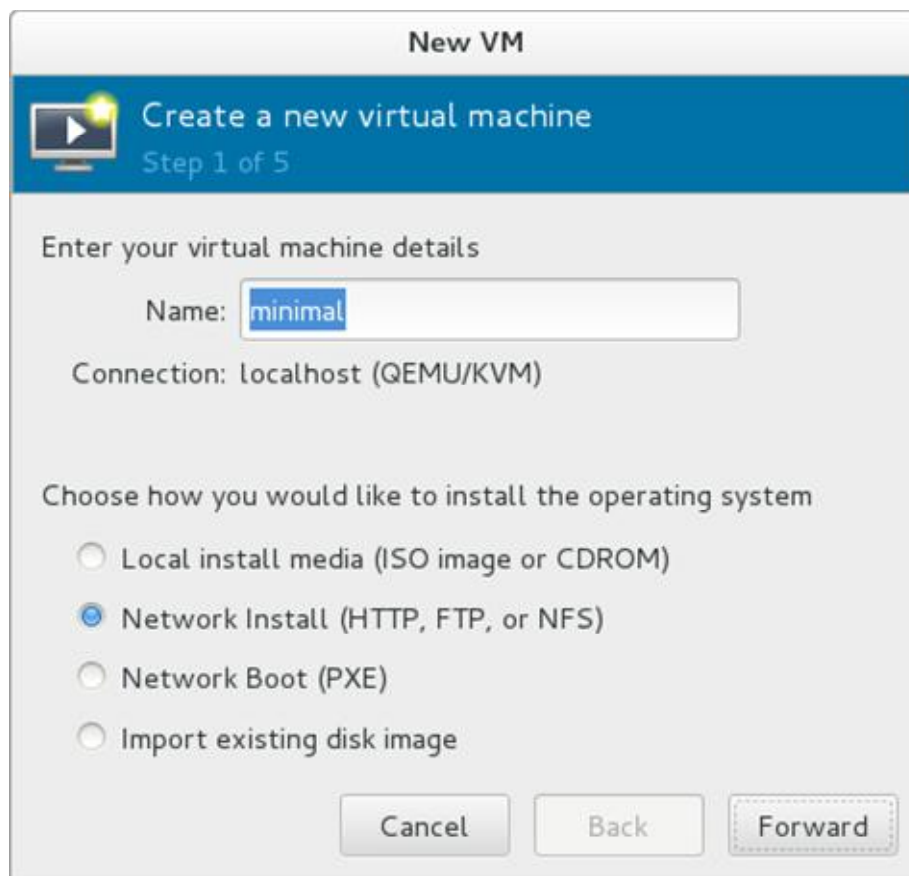



Figure 15.3: Create a virtual machine (step 1 of 5)

Select a name for the virtual machine, which is used as the configuration domain name. The system host name is configured later during the operating system installation. Choices for installation type depend on what resources are prepared. An operating system DVD must be available on the virtual host or on physical media before selecting **Local install media**, for example, or a network installation server must be available when selecting **Network Boot (PXE)**. Sharing and accessing the installation media using any file sharing protocol is chosen here.



**New VM**

Create a new virtual machine  
Step 2 of 5

Provide the operating system install URL

URL:

▼ URL Options

Kickstart URL:

Kernel options:


Automatically detect operating system based on install media

OS type: Linux

Version: Red Hat Enterprise Linux 7

*Figure 15.4: Create a virtual machine (step 2 of 5)*

Enter the **URL** for the installation media network resource. Configure an appropriate **OS type** and **version** for this guest. When the installation media is currently accessible, selecting **Automatically detect operating system** fills these fields. If the fields remain blank or are incorrect, fix the installation media before continuing. The **Kickstart** options allow unattended installations; a Kickstart server and preconfigured client configuration file are required.



*Figure 15.5: Create a virtual machine (step 3 of 5)*

Configure an appropriate amount of **Memory** and **CPUs**. Red Hat Enterprise Linux 7 requires a minimum of one CPU and 1024MB of memory.

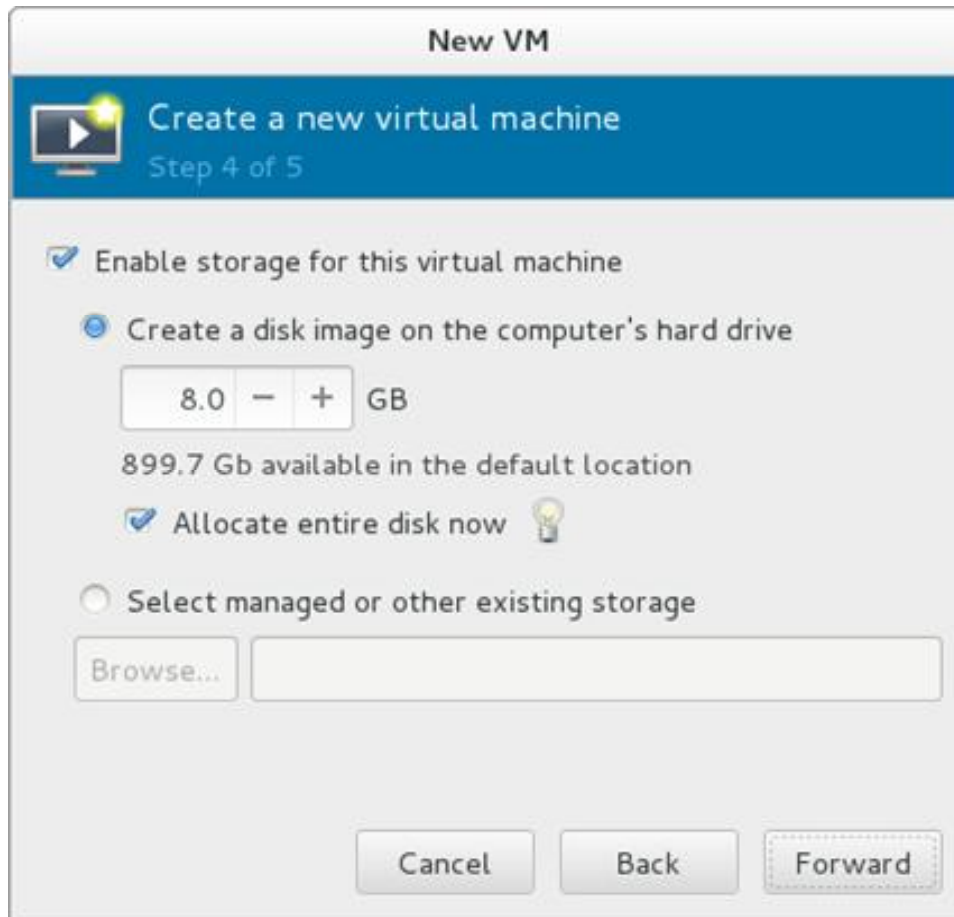


Figure 15.6: Create a virtual machine (step 4 of 5)

Assign storage to the guest virtual machine. For a created disk image, allocating now will provide a small performance benefit, while allocating later saves disk space not yet needed. To **select managed or other existing storage** requires having previously prepared a disk image, physical device, or logical volume prior to beginning this install.

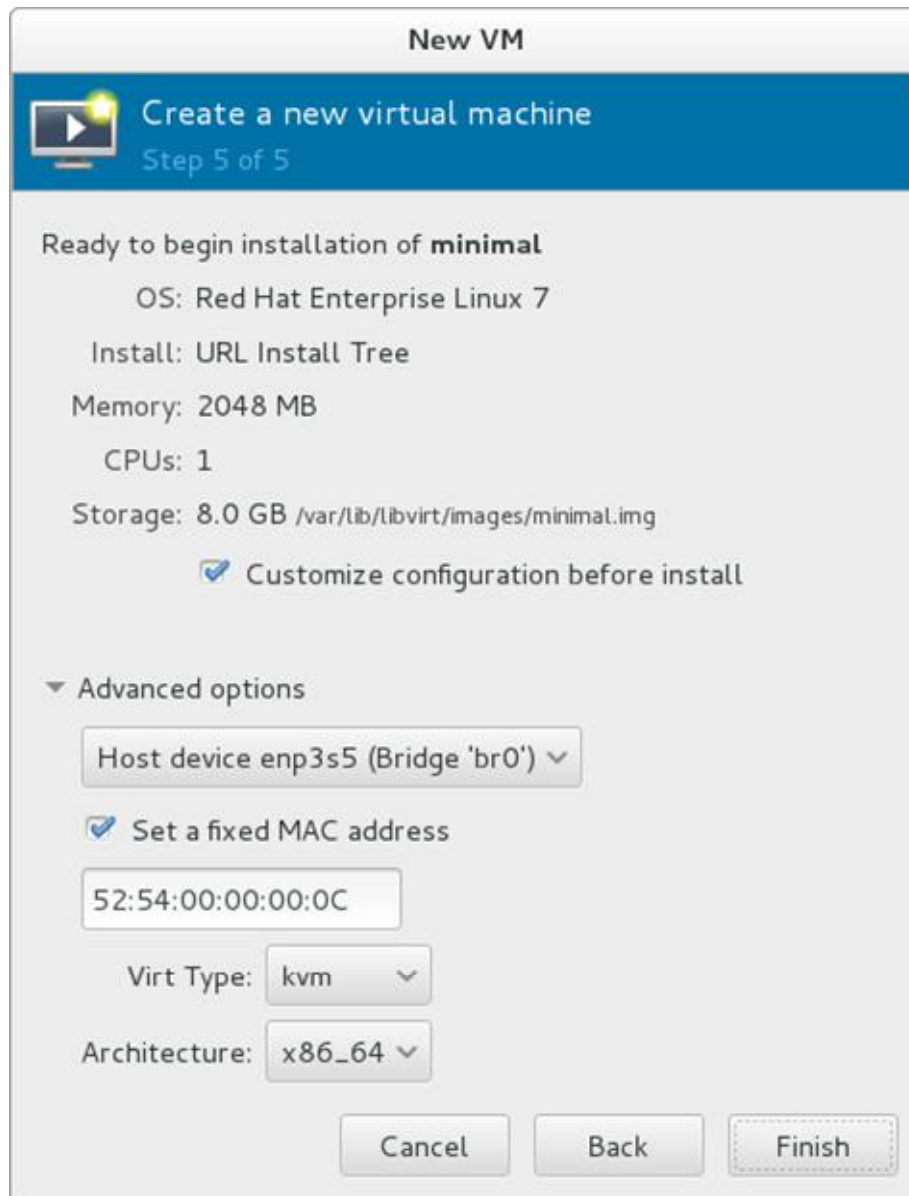


Figure 15.7: Create a virtual machine (step 5 of 5)

Verify the settings of the virtual machine in the top half of the dialog window. Open the **Advanced options** to configure networking. Virtual machines can be configured as **private guests using Network Address Translation (NAT)**, or given the appearance of direct subnet access by using the virtual host's **preconfigured bridge**. **操作系统使用的MAC地址范围**

The **New VM** wizard automatically generates a **MAC address in the 52:54:00 range for virtual guests**. Use the random MAC provided or replace the address with one that fits the environment requirements. Unless the **Customize configuration before install** checkbox is selected, the virtual machine creation will begin when the **Finish** button is pressed.

#### Note

In production environments, DNS and DHCP databases are preconfigured with host names and reserved IP addresses assigned to registered MAC addresses so as to centralize and coordinate enterprise-wide virtual system management. In a Red Hat Training environment, the classroom DNS server has been populated with MAC addresses and unique names for anticipated student guest systems. Use an appropriate MAC address as directed by the instructor or exercise instructions.

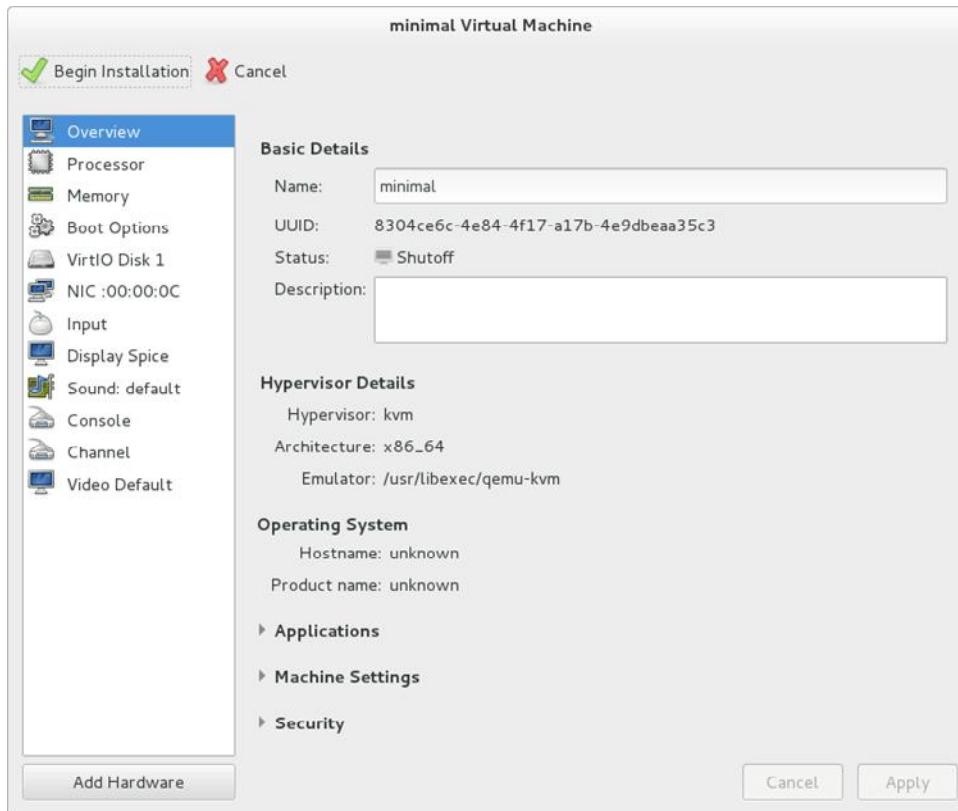


Figure 15.8: Create a virtual machine, customize configuration

If the **Customize configuration before install** checkbox was selected on the last step of the **New VM** wizard, the **domain\_name Virtual Machine** detail screen displays. Administrators can further modify or correct the configuration before the installation begins. Additional hardware—such as graphics and other add-in cards, disks, and network interfaces—may be configured here. When customization is complete, press **Begin Installation** in the upper-left corner. When the installation is complete, administrators can return to this same detail screen to modify or correct the virtual machine configuration. The option for automatically starting this virtual machine when the physical host boots is also found here.

## ▮ Installing Red Hat Enterprise Linux

It is recommended that Red Hat Enterprise Linux be installed using the graphical interface, known as **anaconda**. If anaconda detects that an installation started in text mode where a VNC connection might be possible, anaconda asks to verify. The text mode installation is simpler, but certain options available in graphical mode are not available in text mode.

During the installation, virtual consoles accessed from the physical machine provide information such as diagnostic messages and the ability to enter commands from a shell prompt. This capability is not available during a remote installation. The following table lists the virtual consoles available during installation, and the keystrokes used to switch between them.

▮

### Description of virtual consoles

| Console | Keyboard shortcut | Contents |
|---------|-------------------|----------|
|---------|-------------------|----------|



| Console | Keyboard shortcut  | Contents                                                                                                                     |
|---------|--------------------|------------------------------------------------------------------------------------------------------------------------------|
| 1       | <b>Ctrl+Alt+F1</b> | Main installer console, contains debugging information from <b>anaconda</b>                                                  |
| 2       | <b>Ctrl+Alt+F2</b> | Shell prompt with root access, contains a stored history of useful commands                                                  |
| 3       | <b>Ctrl+Alt+F3</b> | Installation log, displays messages stored in <b>/tmp/anaconda.log</b>                                                       |
| 4       | <b>Ctrl+Alt+F4</b> | Storage log, displays messages on related storage devices from kernel and system services, stored in <b>/tmp/storage.log</b> |
| 5       | <b>Ctrl+Alt+F5</b> | Program log, displays messages from other system utilities, stored in <b>/tmp/program.log</b>                                |
| 6       | <b>Ctrl+Alt+F6</b> | Spare shell prompt                                                                                                           |
| 7       | <b>Ctrl+Alt+F7</b> | The default console with GUI                                                                                                 |

At the **Welcome to Red Hat Enterprise Linux 7.0** screen, select the language to be used during installation. After install, users select their own preferred language upon new account login.

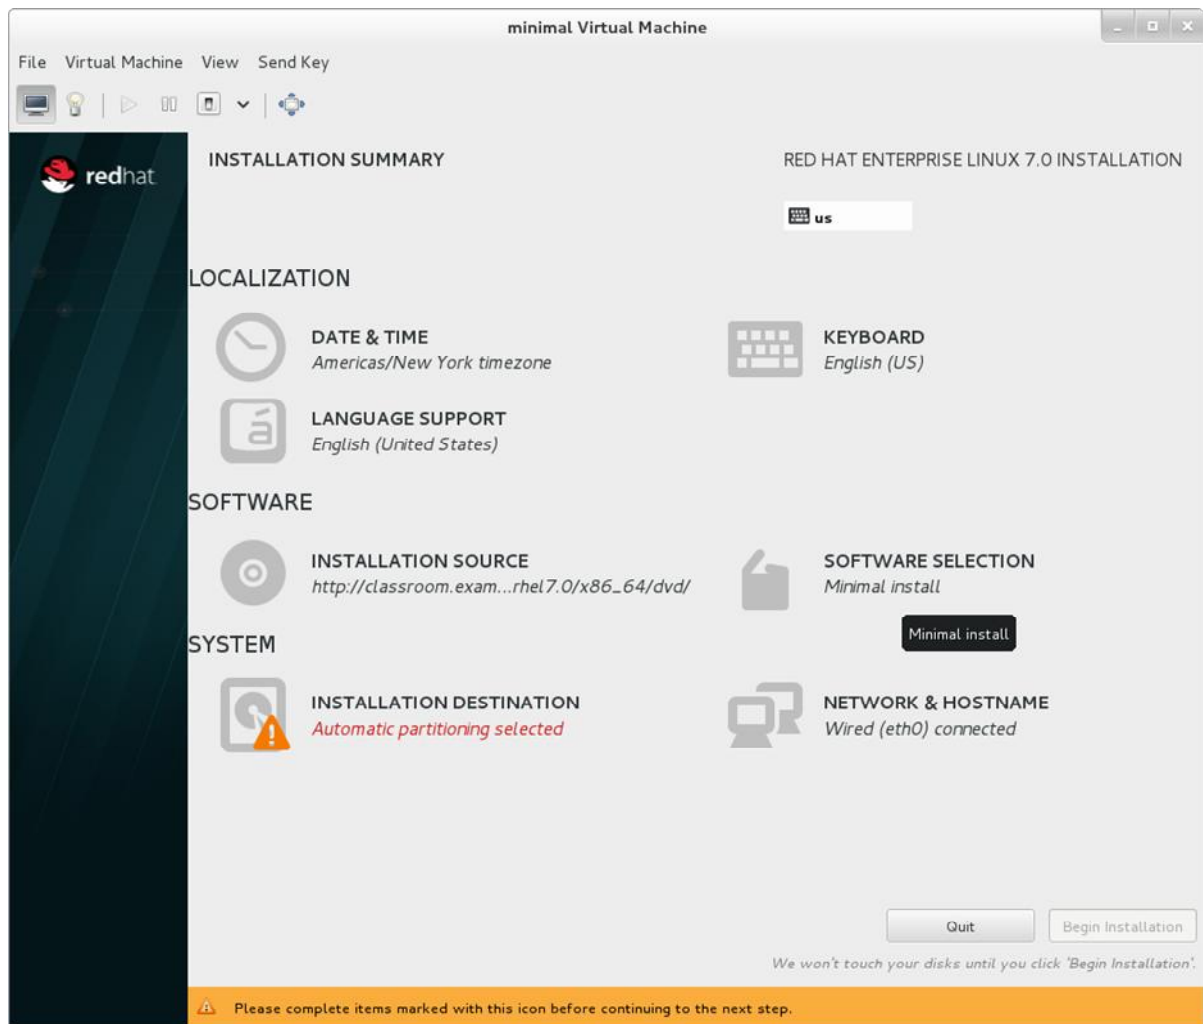


Figure 15.9: Anaconda installer, installation summary

The **anaconda** **Installation Summary** screen displays. From this central screen, installation customizations may be provided before beginning the installation. The installer allows configuring the installation items in any order. Select an item to view or edit. When an item is completed, or it is intended to be completed later, press **Done** to return to this central screen.

Only items marked with a warning symbol are mandatory. The orange status bar at the bottom of the screen warns that these items must be completed before the installation can begin. Once required items are complete, press the **Begin Installation** button. Pressing the **Quit** button aborts the installation. The virtual machine configuration will remain, but be unbootable until an operating system installation is restarted and completed.

As required, complete the following items:

- **Date and time** - Select your city by clicking in the interactive map or select from the dropdown list. Specify a time zone even when using Network Time Protocol (NTP). Messages complaining that NTP is not configured may be ignored if no NTP servers are available.
- **Language support** - Select languages to install in addition to the default language already specified. Multiple languages and locales may be selected.
- **Keyboard** - This item allows adding additional keyboard layouts beyond that included with your default language.

- **Installation source** - The installation source was selected during the virtual machine creation steps and should not need to be modified.
- **Software selection** - By default, the graphical installer selects the Minimal install environment, providing only packages essential to run Red Hat Enterprise Linux. In Software Selection, choose from a list of other base environments by clicking the radio button for an available environment listed in the left pane. Next, select Add-Ons for the selected environment from the checkboxes in the right pane.
- **Installation destination** - Select and partition the disks onto which Red Hat Enterprise Linux will install. This item expects an administrator to comprehend partitioning schemes and file system selection criteria. The default radio button for automatic partitioning will allocate the selected storage devices using all available space. If there are no other operating systems already installed on this computer or you have chosen not to preserve a previously installed operating systems, **anaconda** automatically installs GRUB2 as the boot loader.
- **Network & host name** - Detected network connections are listed in the left pane. Click a listed connection to display more details. To configure a network connection manually, click the **Configure** button in the lower-right corner. A dialog appears to configure the selected connection. Configuration options depend on available network hardware. Only network connections required for installation must be configured here.

When installation customization is complete, press **Begin Installation**.

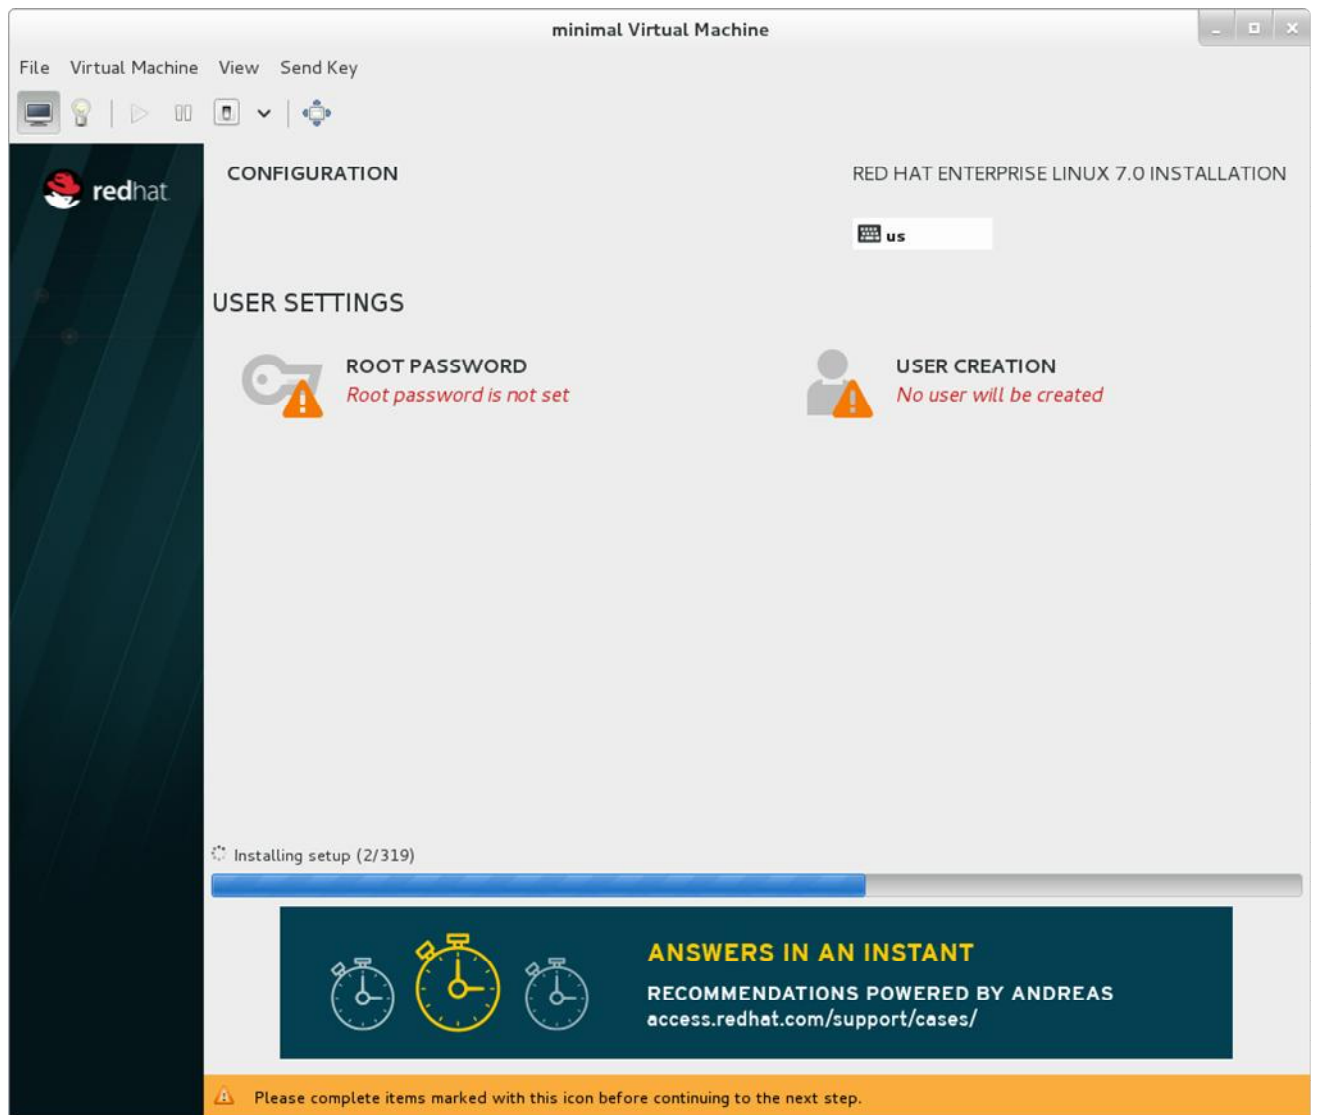


Figure 15.10: Anaconda installer, configuration

The **anaconda** **Configuration** screen displays. The installation has already started; Red Hat Enterprise Linux reports progress near the bottom of the screen as it installs selected packages to the system.

As required, complete the following items:

- **Root password** - The installation program prompts to set a root password for the system. The final stage of the installation process will not continue until a root password is entered.
- **User creation** - Use the root account only for system administration. Create a non-root account for general use. Though recommended to do during installation, this step is optional and can be performed after the installation is complete.

When the installation states that it is complete, press **Reboot**. After the virtual machine's normal power-up sequence has completed, Red Hat Enterprise Linux loads and starts, hidden behind a graphical screen that displays a progress bar. If a graphical desktop was installed, the GUI login displays. Log in as the user created during installation, or as root. The **anaconda Initial Setup** screen displays.

As required, complete the following items:

- **License information** - The installation program prompts to accept the licensing terms.

- **User creation** - Displays only if a non-root account has not yet been created.

When initial setup tasks are complete, press **Finish Installation**. The **FirstBoot** utility displays, prompting entry of the final configuration information for this system. Use this utility to:

- Configure the Kdump crash dump mechanism.
- Set the system time and date.
- Register the machine with Red Hat Network through Subscription Management.

The Red Hat Enterprise Linux installation is now complete. The system is ready for normal login and use.

**References**

Additional information may be available in the section on creating guests with virt-manager in the *Red Hat Enterprise Linux Virtualization Deployment and Administration Guide* for Red Hat Enterprise Linux 7, which can be found at <http://docs.redhat.com/>

**virt-manager**(1) man page

## 15.4. Practice: Installing a New Virtual Machine

### Quiz

The steps to install Red Hat Enterprise Linux using the **anaconda** graphical interface are shown below. Indicate the order the steps should be taken.

- Name the virtual machine and select an installation source.
- Provide the source location and operating system type
- Enter the Memory and CPU settings
- Configure storage for this virtual machine
- Configure networking options for the installation
- Modify Localization parameters for Date and Time, Language and Keyboard
- Select the Software to be installed
- Provide System parameters for disk partitioning, networking and hostname
- Provide User Settings for a root password and a non-root account

## 15.5. Chapter Test: Using Virtualized Systems

### Quiz

Match the following items to their counterparts in the table.

| Configuration description                                 | Red Hat product                   |
|-----------------------------------------------------------|-----------------------------------|
| Single system hardware providing KVM support              | Red Hat Enterprise Linux          |
| Multiple system hardware providing virtualized redundancy | Red Hat Enterprise Virtualization |
| Multiple system hardware providing private cloud          | RHEL OpenStack platform           |
| Cloud provider providing public cloud                     | OpenStack in public cloud         |
| Management utility for standalone KVM hosts               | virt-manager                      |
| Management utility for multiple host virtualization       | RHEV-M                            |

| Configuration description                                              | Red Hat product |
|------------------------------------------------------------------------|-----------------|
| platform                                                               |                 |
| Management utility for all virtualization and cloud platforms combined | CloudForms      |

# Chapter 16. Comprehensive Review

## 16.1. Red Hat System Administration I Comprehensive Review

### 16.2. Lab: Comprehensive Review

#### Introduction

| Overview   |                                                                                                               |
|------------|---------------------------------------------------------------------------------------------------------------|
| Goal       | To practice and demonstrate knowledge and skills learned in Red Hat System Administration I.                  |
| Objectives | <ul style="list-style-type: none"> <li>Practice skills learned in Red Hat System Administration I.</li> </ul> |
| Lab        | <ul style="list-style-type: none"> <li>Comprehensive Review</li> </ul>                                        |

## 16.1. Red Hat System Administration I Comprehensive

### Review

#### Objectives

After completing this section, students should be able to demonstrate knowledge and skill of the topics covered in each chapter.

#### Reviewing Red Hat System Administration I

Before beginning the comprehensive review for this course, students should be comfortable with the topics covered in each chapter.

Students can refer to earlier sections in the textbook for extra study.

#### Chapter 1, *Accessing the Command Line*

To log into a Linux system and run simple commands using the shell.

- Use Bash shell syntax to enter commands at a Linux console.
- Launch applications in a GNOME desktop environment.
- Use Bash features to run commands from a shell prompt using fewer keystrokes.

#### Chapter 2, *Managing Files From the Command Line*

To copy, move, create, delete, and organize files while working from the Bash shell prompt.

- Identify the purpose for important directories on a Linux system.
- Specify files using absolute and relative path names.

- Create, copy, move, and remove files and directories using command-line utilities.
- Match one or more file names using shell expansion as arguments to shell commands.

### **Chapter 3, *Getting Help in Red Hat Enterprise Linux***

To resolve problems by using on-line help systems and Red Hat support utilities.

- Use the man Linux manual reader.
- Use the pinfo GNU Info reader.
- Use the Red Hat Package Manager (RPM) package documentation.
- Use the redhat-support-tool command.

### **Chapter 4, *Creating, Viewing, and Editing Text Files***

To create, view, and edit text files from command output or in an editor.

- Redirect the text output of a program to a file or to another program.
- Edit existing text files and create new files from the shell prompt with a text editor.
- Copy text from a graphical window to a text file using a text editor running in the graphical environment.

### **Chapter 5, *Managing Local Linux Users and Groups***

To manage local Linux users and groups and administer local password policies.

- Explain the role of users and groups on a Linux system and how they are understood by the computer.
- Run commands as the superuser to administer a Linux system.
- Create, modify, lock, and delete locally defined user accounts.
- Create, modify, and delete locally defined group accounts.
- Lock accounts manually or by setting a password-aging policy in the shadow password file.

### **Chapter 6, *Controlling Access to Files with Linux File System Permissions***

To set Linux file system permissions on files and interpret the security effects of different permission settings.

- Explain how the Linux file permissions model works.
- Change the permissions and ownership of files using command-line tools.
- Configure a directory in which newly created files are automatically writable by members of the group which owns the directory, using special permissions and default umask settings.

### **Chapter 7, *Monitoring and Managing Linux Processes***

To evaluate and control processes running on a system.

- List and interpret basic information about processes running on the system.
- Control processes in the shell's session using bash job control.
- Terminate and control processes using signals.
- Monitor resource usage and system load due to process activity.

### **Chapter 8, *Controlling Services and Daemons***

To control and monitor network services and system daemons using systemd.

- List system daemons and network services started by the systemd service and socket units.
- Control system daemons and network services using **systemctl**.

### **Chapter 9, *Configuring and Securing OpenSSH Service***

To configure secure command-line access on remote systems using OpenSSH.

- Log into a remote system using `ssh` to run commands from a shell prompt.
- Set up `ssh` to allow secure password-free logins by using a private authentication key file.
- Customize `sshd` configuration to restrict direct logins as root or to disable password-based authentication.

### **Chapter 10, *Analyzing and Storing Logs***

To locate and accurately interpret relevant system log files for troubleshooting purposes.

- Describe the basic syslog architecture in Red Hat Enterprise Linux 7.
- Interpret entries in relevant syslog files to troubleshoot problems or review system status.
- Find and interpret log entries in the `systemd` journal to troubleshoot problems or review system status.
- Configure `systemd-journald` to store its journal on disk rather than in memory.
- Maintain accurate time synchronization and time zone configuration to ensure correct timestamps in system logs.

### **Chapter 11, *Managing Red Hat Enterprise Linux Networking***

To configure basic IPv4 networking on Red Hat Enterprise Linux systems.

- Explain fundamental concepts of computer networking.
- Test and review current network configuration with basic utilities.
- Manage network settings and devices with `nmcli` and NetworkManager.
- Modify network settings by editing the configuration files.
- Configure and test system host name and name resolution.

### **Chapter 12, *Archiving and Copying Files Between Systems***

To archive and copy files from one system to another.

- Use `tar` to create new compressed archive files and extract files from existing archive files.
- Copy files securely to or from a remote system running `sshd`.
- Securely synchronize the contents of a local file or directory with a remote copy.

### **Chapter 13, *Installing and Updating Software Packages***

To download, install, update, and manage software packages from Red Hat and YUM package repositories.

- Register systems with your Red Hat account and entitle them to software updates for installed products.
- Explain what an RPM package is and how RPM packages are used to manage software on a Red Hat Enterprise Linux system.
- Find, install, and update software packages using the `yum` command.
- Enable and disable use of Red Hat or third-party YUM repositories.
- Examine and install downloaded software package files.

### **Chapter 14, *Accessing Linux File Systems***

To access and inspect existing file systems on a Red Hat Enterprise Linux system.



- Identify the file system hierarchy.
- Access the contents of file systems.
- Use hard links and symlinks to make multiple names.
- Search for files on mounted file systems.

### Chapter 15, *Using Virtualized Systems*

To create and use Red Hat Enterprise Linux virtual machines with Kernel-based Virtual Machine (KVM) and libvirt.

- Install a Red Hat Enterprise Linux system as a host for running virtual machines.
- Perform an interactive install of Red Hat Enterprise Linux on a virtual machine.

### References

Get information on more classes available from Red Hat at <http://www.redhat.com/training/>

## 16.2. Lab: Comprehensive Review

### Performance checklist

In this lab, you will practice and demonstrate your knowledge and skills.

#### Outcomes:

Complete the following tasks and successfully grade the serverX system with **lab sal-review grade** as user root on serverX.

#### Before you begin...

Reset the serverX machine.

Run the **lab sal-review setup** as user root on serverX.

1. Use Bash commands to complete the following tasks on the serverX machine:
  - Display the first 12 lines of the **/usr/bin/clean-binary-files** file and send the output to the **/home/student/headtail.txt** file.
  - Display the last nine lines of the **/usr/bin/clean-binary-files** file and add the output to the **/home/student/headtail.txt** file.

Show Solution

2. Ten new Linux systems require change documentation files. Carry out the following tasks on serverX to create them:
  - Create the empty files with the file name **system\_changes-machineY-month\_Z.txt** in the **/home/student** directory on the serverX machine as user student. Replace Y with the machine number and replace Z with the months *jan*, *feb*, and *mar*.
  - Create the **/home/student/syschanges** directory with the subdirectories **jan**, **feb**, and **mar**.
  - Sort all newly created files by month into the corresponding subdirectory.
  - Remove all newly created files related to machine 9 and 10, because the hardware has been replaced permanently.

Show Solution

3. Use the man pages to research how to turn off the use of colors in the output. Put the relevant option of the **ls** command into the text file **/home/student/lscolor.txt** on serverX.

Show Solution

4. Copy the file **/home/student/vimfile.txt** to **/home/student/longlisting.txt** on serverX. Use the **vim** editor to change the **/home/student/longlisting.txt** file according to the following requirements:

- Remove the file owner column. Do not remove any spaces.
- Remove the **Documents** and **Pictures** rows.
- Save the file when done with editing.

Show Solution

5. Change configuration and add new users and a new group according to the following requirements:
  - Change the default system settings for newly created users to ensure their passwords are changed at least every 60 days.
  - Create a new group named **instructors** with a GID of 30000.
  - Create three new users: **gorwell**, **rbradbury**, and **dadams**, with a password of **firstpw**.
  - Add the new users to the supplementary group **instructors**. The primary group should remain as the user private group.
  - Set the three newly created accounts to expire 60 days from today.
  - Change the password policy for the **gorwell** account to require a new password every 10 days.
  - Force all three newly created users to change their password on first login.

Show Solution

6. Create the shared directory **/home/instructors** on serverX according to the following requirements:
  - The directory is owned by user **root** and group **instructors**.
  - Set permissions on the **/home/instructors** directory so it has the GID bit set on the directory, the owner and group have full read/write/execute permissions, and other users have read permission to the directory.

Show Solution

7. Determine the process using the most CPU resources on serverX and terminate it.

Show Solution

8. Stop the currently running cups printing service on serverX. The service should not get automatically started on system boot.

Show Solution

9. Configure the ssh service on serverX according to the following requirements:
  - User student on serverX can log in with a SSH public key to the student account on desktopX.
  - Disable **ssh** login for the root user and password-based SSH authentication on serverX.

Show Solution

10. Your serverX machine has been relocated to the Bahamas. The following changes have to be made on the serverX machine:
  - Change the time zone on the serverX machine to Bahamas and verify the time zone has been changed properly.

Show Solution

11. Record the command to display all **systemd** journal entries recorded between 9:05:00 and 9:15:00 in the **/home/student/systemdreview.txt** file.

Show Solution

12. Configure **rsyslogd** by adding a rule to the newly created configuration file **/etc/rsyslog.d/auth-errors.conf** to log all security and authentication messages that get recorded in the authpriv facility with the priority alert and higher to the **/var/log/auth-errors** file as well. Test the newly added log directive with the **logger** command.

Show Solution

13. Create a new static network connection using the settings in the following table. Be sure to replace the **x** with the correct number for your systems.
  - Configure the new connection to be automatically started.
  - Other connections should not start automatically.
  - Modify the new connection so that it also uses the address 10.0.X.1/24.
  - Configure the **hosts** file so that 10.0.X.1 can be referenced as "myhost".
  - Set the host name to serverX.example.com.

| Parameter       | Setting           |
|-----------------|-------------------|
| Connection name | review            |
| IP address      | 172. 25. x. 11/16 |
| Gateway address | 172. 25. 0. 1     |
| DNS address     | 172. 25. 254. 254 |

Show Solution

14. Synchronize the `/etc` directory tree on serverX to the `/configbackup` directory on serverX.

Show Solution

15. Create an archive named `/root/configuration-backup-server.tar.gz` with the `/configbackup` directory as content.

Show Solution

16. To prepare the archived directory tree for comparison with the currently actively used configuration files on serverX, extract the contents of the `/root/configuration-backup-server.tar.gz` archive to the `/tmp/configcompare/` directory on serverX.

Show Solution

17. Perform the following tasks on the serverX machine:

- Use `ssh` to execute the `hostname` command on desktopX as user student. Send the output of the `hostname` command to the `/tmp/scpfile.txt` file on desktopX.
- Use `scp` to copy the `/tmp/scpfile.txt` file from desktopX to `/home/student/scpfile.txt`.

Show Solution

18. Create the file `/etc/yum.repos.d/localupdates.repo` to enable the “Updates” repository found on the content machine. It should access content found at the following URL: `http://content.example.com/rhel7.0/x86_64/errata`. Do *not* check GPG signatures.

Show Solution

19. Configure serverX to adhere to specific software requirements:

- The `kernel` package should be updated to the latest version.
- The `xsane-gimp` package should be installed.
- The `rht-system` package should be installed.
- For security reasons, serverX should not have the `wvdial` package installed.

Show Solution

20. Generate a disk usage report with the `du` command of the `/usr/share/fonts` directory on serverX and save the result in the `/home/student/dureport.txt` file.

Show Solution

21. Identify and mount a newly added file system by UUID on the `/mnt/datadump` directory on serverX.

Show Solution

22. Create the soft link `/root/mydataspace`, which points to the `/mnt/datadump` directory on serverX.

Show Solution

23. Record the command to find all soft links on serverX that have `data` as part of their name in the `/home/student/find.txt` file.

Show Solution

Reset the serverX machine.

Run the `lab sal-review setup` as user `root` on serverX.

1. Use Bash commands to complete the following tasks on the serverX machine:
  - o Display the first 12 lines of the `/usr/bin/clean-binary-files` file and send the output to the `/home/student/headtail.txt` file.
  - o Display the last nine lines of the `/usr/bin/clean-binary-files` file and add the output to the `/home/student/headtail.txt` file.
    - c. Display the first 12 lines of the `/usr/bin/clean-binary-files` file and send the command output to the `/home/student/headtail.txt` file.

```
[student@serverX ~]$ head -n 12
/usr/bin/clean-binary-files >/home/student/headtail.txt
```

- d. Display the last nine lines of the `/usr/bin/clean-binary-files` file and add the command output to the `/home/student/headtail.txt` file.

```
[student@serverX ~]$ tail -n 9
/usr/bin/clean-binary-files >>/home/student/headtail.txt
```

2. Ten new Linux systems require change documentation files. Carry out the following tasks on serverX to create them:
  - o Create the empty files with the file name `system_changes-machineY-month_Z.txt` in the `/home/student` directory on the serverX machine as user `student`. Replace `Y` with the machine number and replace `Z` with the months `jan`, `feb`, and `mar`.
  - o Create the `/home/student/syschanges` directory with the subdirectories `jan`, `feb`, and `mar`.

- o Sort all newly created files by month into the corresponding subdirectory.
  - o Remove all newly created files related to machine 9 and 10, because the hardware has been replaced permanently.
- d. Create a total of 30 files with names **system\_changes-machineY-month\_Z.txt**. Replace Y with the machine number and replace Z with the months *jan*, *feb*, and *mar*.

```
[student@serverX ~]$ touch
~student/system_changes-machine{1..10}-month_{
jan,feb,mar}.txt
```

- e. Create the **/home/student/syschanges** directory with the subdirectories **jan**, **feb**, and **mar**.

```
[student@serverX ~]$ mkdir -p
/home/student/syschanges/{jan,feb,mar}
```

- f. Sort all newly created files by month into the corresponding subdirectory.

```
[student@serverX ~]$ mv
~student/system_changes-machine*jan.txt
/home/student/syschanges/jan/
[student@serverX ~]$ mv
~student/system_changes-machine*feb.txt
/home/student/syschanges/feb/
```

```
[student@serverX ~]$ mv
~student/system_changes-machine*mar.txt
/home/student/syschanges/mar/
```

- g. Remove all newly created files related to machine 9 and 10.

```
[student@serverX ~]$ rm -f
/home/student/syschanges/*/system_changes-mach
ine{9,10}*.txt
```

3. Use the man pages to research how to turn off the use of colors in the output. Put the relevant option of the **ls** command into the text file **/home/student/lscolor.txt** on serverX.

- . Look up the relevant option in the **ls(1)** man page to determine how to prevent ls from providing colorful output. What is the correct option?

```
[student@serverX ~]$ man ls
```

**ls** uses **--color=never** to turn off colors in the command output.

- a. Create the text file `/home/student/lscolor.txt` with the `ls` option to turn off colorful output.

```
[student@serverX ~]$ echo
"--color=never" >/home/student/lscolor.txt
```

4. Copy the file `/home/student/vimfile.txt` to `/home/student/longlisting.txt` on serverX. Use the `vim` editor to change the `/home/student/longlisting.txt` file according to the following requirements:

- Remove the file owner column. Do not remove any spaces.
- Remove the **Documents** and **Pictures** rows.
- Save the file when done with editing.

- c. Copy the file `/home/student/vimfile.txt` to `/home/student/longlisting.txt`.

```
[student@serverX ~]$ cp
/home/student/vimfile.txt
/home/student/longlisting.txt
```

- d. Edit the file using Vim, to take advantage of *visual mode*.

```
[student@serverX ~]$ vim
/home/student/longlisting.txt
```

- e. Remove the *owner* column from the file.

Use the arrow keys to position the cursor at the first character of the group owner column. Enter visual mode with **Ctrl-v**. Use the arrow keys to position the cursor at the last character and row of the group owner column. Delete the selection with **x**.

- f. Remove the **Documents** and **Pictures** rows. This time, enter visual mode with an uppercase **v**, which automatically selects full lines.

Use the arrow keys to position the cursor at any character on the **Documents** row. Enter visual mode with an uppercase **V**. The full line is selected, as shown in the screen shot. Delete the selection with **x**. Repeat for the **Pictures** row.

- g. Save the file and exit the editor.

Press the "esc" key and enter `":wq"` to write the file and exit `vim`.

5. Change configuration and add new users and a new group according to the following requirements:

- Change the default system settings for newly created users to ensure their passwords are changed at least every 60 days.
  - Create a new group named **instructors** with a GID of 30000.
  - Create three new users: **gorwell**, **rbradbury**, and **dadams**, with a password of **firstpw**.
  - Add the new users to the supplementary group **instructors**. The primary group should remain as the user private group.
  - Set the three newly created accounts to expire 60 days from today.
  - Change the password policy for the **gorwell** account to require a new password every 10 days.
  - Force all three newly created users to change their password on first login.
- g. Change the default system settings for newly created users to ensure their passwords are changed at least every 60 days.

```

h. [student@serverX ~]$ sudo vim /etc/login.defs
i. [student@serverX ~]$ cat /etc/login.defs
j. ...Output omitted...
k. PASS_MAX_DAYS 60
l. PASS_MIN_DAYS 0
m. PASS_MIN_LEN 5
n. PASS_WARN_AGE 7
o. ...Output omitted...

```

- p. Create a new group named **instructors** with a GID of 30000.

```

q. [student@serverX ~]$ sudo groupadd -g 30000
 instructors
r. [student@serverX ~]$ tail -5 /etc/group
s. stapdev:x:158:
t. pesign:x:989:
u. tcpdump:x:72:
v. slocate:x:21:
w. instructors:x:30000:

```

- x. Create three new users: **gorwell**, **rbradbury**, and **dadams**, with a password of **firstpw** and add them to the supplementary group **instructors**. The primary group should remain as the user private group.

```

y. [student@serverX ~]$ sudo useradd -G
 instructors gorwell

```



```

z. [student@serverX ~]$ sudo useradd -G
 instructors rbradbury
aa. [student@serverX ~]$ sudo useradd -G
 instructors dadams
bb. [student@serverX ~]$ tail -5 /etc/group
cc. slocate:x:21:
dd. instructors:x:30000:gorwell,rbradbury,dadams
ee. gorwell:x:1001:
ff. rbradbury:x:1002:
gg. dadams:x:1003:
hh. [student@serverX ~]$ sudo passwd gorwell
ii. Changing password for user gorwell.
jj. New password: firstpw
kk. BAD PASSWORD: The password is shorter than 8 characters
ll. Retype new password: firstpw
mm. passwd: all authentication tokens updated successfully.
nn. [student@serverX ~]$ sudo passwd rbradbury
oo. [student@serverX ~]$ sudo passwd dadams

```

pp. Determine the date 60 days in the future and set each of the three new user accounts to expire on that date.

```

qq. [student@serverX ~]$ date -d "+60 days"
rr. Mon April 5 11:49:24 EDT 2014
ss. [student@serverX ~]$ sudo chage -E 2014-04-05
 gorwell
tt. [student@serverX ~]$ sudo chage -E 2014-04-05
 rbradbury
uu. [student@serverX ~]$ sudo chage -E 2014-04-05
 dadams

```

vv. Change the password policy for the **gorwell** account to require a new password every 10 days.

```

ww. [student@serverX ~]$ sudo chage -M 10 gorwell
xx. [student@serverX ~]$ chage -l gorwell
yy. Last password change : Feb 04, 2014
zz. Password expires : Feb 14, 2014
aaa. Password inactive : never

```

```

bbb. Account expires : April
 05, 2014

ccc. Minimum number of days between password change : 0

ddd. Maximum number of days between password change : 10

eee. Number of days of warning before password expires : 7

```

fff. Force all three newly created users to change their password on first login.

```

ggg. [student@serverX ~]$ sudo chage -d 0
 gorwell

hhh. [student@serverX ~]$ sudo chage -d 0
 rbradbury

iii. [student@serverX ~]$ sudo chage -d 0
 dadams

```

6. Create the shared directory **/home/instructors** on serverX according to the following requirements:
  - o The directory is owned by user root and group instructors.
  - o Set permissions on the **/home/instructors** directory so it has the GID bit set on the directory, the owner and group have full read/write/execute permissions, and other users have read permission to the directory.

b. Open a terminal window and become root on serverX.

```

[student@serverX ~]$ su -
Password: redhat

```

```

[root@serverX ~]#

```

c. Create the **/home/instructors** directory.

```

[root@serverX ~]# mkdir /home/instructors

```

d. Change group permissions on the **/home/instructors** directory so it belongs to the group instructors.

```

[root@serverX ~]# chown :instructors
/home/instructors

```

e. Set permissions on the **/home/instructors** directory so it is a set GID bit directory (2), the owner (7) and group (7) have full read/write/execute permissions, and other users have read permission (4) to the directory.

```

[root@serverX ~]# chmod 2774 /home/instructors

```

f. Check that the permissions were set properly.

```
[root@serverX ~]# ls -ld /home/instructors
```

```
drwxrwsr--. 2 root instructors 1024 Dec 9 1:38 /home/instructors
```

7. Determine the process using the most CPU resources on serverX and terminate it.

. In a terminal window, run the **top** utility. Size the window as tall as possible. Top sorts all processes by CPU utilization. The cpuhog process is the one with the highest CPU usage.

```
a. [root@serverX ~]# top
b. top - 12:47:46 up 2:02, 3 users, load average: 1.67, 1.25, 0.73
c. Tasks: 361 total, 6 running, 355 sleeping, 0 stopped, 0 zombie
d. %Cpu(s): 98.5 us, 1.4 sy, 0.0 ni, 0.0 id, 0.0 wa, 0.0 hi, 0.1 si,
 0.0 st
e. KiB Mem: 2043424 total, 897112 used, 1146312 free, 1740 buffers
f. KiB Swap: 4079612 total, 0 used, 4079612 free. 296276 cached
 Me
g.
h. PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND
i. 4019 root 20 0 4156 76 0 R 57.5 0.0 2:54.15 cpuhog
j. 2492 student 20 0 1359500 168420 37492 S 16.8 8.2 3:55.58
 gnome-shell
k. 1938 root 20 0 189648 35972 7568 R 1.9 1.8 0:29.66 Xorg
l. 2761 student 20 0 620192 19688 12296 S 0.4 1.0 0:04.48
 gnome-termi+
```

```
output truncated
```

m. Exit the **top** display.

Press **q** to quit.

n. Terminate the cpuhog process using the command line. Confirm that the processes no longer display in **top**.

```
[root@serverX ~]# pkill cpuhog
```

8. Stop the currently running cups printing service on serverX. The service should not get automatically started on system boot.

. Stop the **cups** service.

```
[student@serverX ~]$ sudo systemctl stop cups
```

```
[student@serverX ~]$ sudo systemctl status cups
```

- a. Configure the **cups** service so that it does not start at system boot.

```
[student@serverX ~]$ sudo systemctl disable cups
```

```
[student@serverX ~]$ sudo systemctl status cups
```

9. Configure the ssh service on serverX according to the following requirements:

- o User student on serverX can log in with a SSH public key to the student account on desktopX.
- o Disable **ssh** login for the root user and password-based SSH authentication on serverX.

- b. Generate a SSH public key on serverX as user student.

```
c. [student@serverX ~]$ ssh-keygen
```

- d. Install the SSH public key generated previously on serverX to the **student** account on desktopX.

```
e. [student@serverX ~]$ ssh-copy-id desktopX
```

```
f. /usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s),
to filter out any that are already installed
```

```
g. /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
```

```
h. student@desktopX's password: student
```

```
i.
```

```
j. Number of key(s) added: 1
```

```
k.
```

```
l. Now try logging into the machine, with: "ssh 'student@desktopX'"
```

```
m. and check to make sure that only the key(s) you wanted were added.
```

```
n.
```

- o. Log in, then change to the root account, on the serverX virtual machine.

```
[student@desktopX ~]$ su -
```

- p. Customize the ssh service on serverX by disabling SSH connections for the user root and only allow key-based login.

Set the necessary configuration file parameters  
in **/etc/ssh/sshd\_config**:

```
PermitRootLogin no
PasswordAuthentication no
```

- q. Restart the sshd service on serverX.

```
[root@serverX ~]# systemctl restart sshd
```

- r. On a different terminal window on desktopX, validate that user root cannot connect to serverX with the **ssh** command. It should fail because we disabled root logins with the ssh service.

```
s. [student@desktopX ~]$ ssh root@serverX
t. Password: redhat
u. Permission denied, please try again.
v. Password: redhat
w. Permission denied, please try again.
x. Password: redhat
y. Permission denied (publickey,gssapi-keyex,gssapi-with-mic,password).
```

10. Your serverX machine has been relocated to the Bahamas. The following changes have to be made on the serverX machine:

- o Change the time zone on the serverX machine to Bahamas and verify the time zone has been changed properly.
  - a. Identify the correct time zone for Bahamas on serverX.

```
b. [root@serverX ~]# tzselect
c. Please identify a location so that time zone rules can be set correctly.
d. Please select a continent or ocean.
e. 1) Africa
f. 2) Americas
g. 3) Antarctica
h. 4) Arctic Ocean
i. 5) Asia
j. 6) Atlantic Ocean
k. 7) Australia
l. 8) Europe
m. 9) Indian Ocean
n. 10) Pacific Ocean
o. 11) none - I want to specify the time zone using the Posix TZ format.
p. ?? 2
q. Please select a country.
r. 1) Anguilla 28) Haiti
s. 2) Antigua & Barbuda 29) Honduras
t. 3) Argentina 30) Jamaica
```

u. 4) Aruba 31) Martinique  
v. 5) Bahamas 32) Mexico  
w. 6) Barbados 33) Montserrat  
x. 7) Belize 34) Nicaragua  
y. 8) Bolivia 35) Panama  
z. 9) Brazil 36) Paraguay  
aa. 10) Canada 37) Peru  
bb. 11) Caribbean Netherlands 38) Puerto Rico  
cc. 12) Cayman Islands 39) St Barthelemy  
dd. 13) Chile 40) St Kitts & Nevis  
ee. 14) Colombia 41) St Lucia  
ff. 15) Costa Rica 42) St Maarten (Dutch part)  
gg. 16) Cuba 43) St Martin (French part)  
hh. 17) Curacao 44) St Pierre & Miquelon  
ii. 18) Dominica 45) St Vincent  
jj. 19) Dominican Republic 46) Suriname  
kk. 20) Ecuador 47) Trinidad & Tobago  
ll. 21) El Salvador 48) Turks & Caicos Is  
mm. 22) French Guiana 49) United States  
nn. 23) Greenland 50) Uruguay  
oo. 24) Grenada 51) Venezuela  
pp. 25) Guadeloupe 52) Virgin Islands (UK)  
qq. 26) Guatemala 53) Virgin Islands (US)  
rr. 27) Guyana

ss. **#? 5**

tt.

uu. The following information has been given:

vv.

ww. Bahamas

xx.

yy. Therefore TZ='America/Nassau' will be used.

zz. Local time is now: Fri Mar 7 09:38:50 EST 2014.

aaa. Universal Time is now: Fri Mar 7 14:38:50 UTC 2014.

bbb. Is the above information OK?

ccc. 1) Yes

ddd. 2) No

```
eee. #? 1
fff.
ggg. You can make this change permanent for yourself by appending the line
hhh. TZ='America/Nassau'; export TZ
iii. to the file '.profile' in your home directory; then log out and log in
 again.
jjj.
kkk. Here is that TZ value again, this time on standard output so that you
lll. can use the /usr/bin/tzselect command in shell scripts:
mmm. America/Nassau
```

nnn. Change the time zone to America/Nassau on serverX.

```
[root@serverX ~]# timedatectl set-timezone
America/Nassau
```

ooo. Verify the time zone has been properly set on serverX.

```
ppp. [root@serverX ~]# timedatectl
qqq. Local time: Wed 2014-04-09 18:21:06 CEST
rrr. Universal time: Wed 2014-04-09 16:21:06 UTC
sss. RTC time: Wed 2014-04-09 16:21:06
ttt. Timezone: America/Nassau (CEST, +0200)
uuu. NTP enabled: yes
vvv. NTP synchronized: no
www. RTC in local TZ: no
xxx. DST active: yes
yyy. Last DST change: DST began at
zzz. Sun 2014-03-30 01:59:59 CET
aaaa. Sun 2014-03-30 03:00:00 CEST
bbbb. Next DST change: DST ends (the clock jumps one hour backwards) at
cccc. Sun 2014-10-26 02:59:59 CEST
```

```
Sun 2014-10-26 02:00:00 CET
```

11. Record the command to display all **systemd** journal entries recorded between 9:05:00 and 9:15:00 in the **/home/student/systemdreview.txt** file.

```
[root@serverX ~]# echo "journalctl --since 9:05:00
--until 9:15:00" >/home/student/systemdreview.txt
```

12. Configure **rsyslogd** by adding a rule to the newly created configuration file **/etc/rsyslog.d/auth-errors.conf** to log all security and authentication messages that get recorded in the **authpriv** facility with the priority **alert** and higher to the **/var/log/auth-errors** file as well. Test the newly added log directive with the **logger** command.

- . Add the directive to log **authpriv.alert** syslog messages to the **/var/log/auth-errors** file in the **/etc/rsyslog.d/auth-errors.conf** configuration file.

```
[root@serverX ~]# echo "authpriv.alert
/var/log/auth-errors" >/etc/rsyslog.d/auth-err
ors.conf
```

- a. Restart the **rsyslog** service on serverX.

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

- b. Monitor the newly created log file **/var/log/auth-errors** on serverX for changes in a different terminal window.

```
[root@serverX ~]# tail -f /var/log/auth-errors
```

- c. Use **logger** to create a new log entry to **/var/log/auth-errors** on serverX.

```
[root@serverX ~]# logger -p authpriv.alert
"Logging test authpriv.alert"
```

- d. Verify the message sent to syslog with the **logger** command appears in the **/var/log/auth-errors** file on serverX in the terminal running **tail -f /var/log/auth-errors**.

```
[root@serverX ~]# tail -f /var/log/auth-errors
```

```
Feb 13 11:21:53 server1 root: Logging test authpriv.alert
```

13. Create a new static network connection using the settings in the following table. Be sure to replace the **x** with the correct number for your systems.
- Configure the new connection to be automatically started.
  - Other connections should not start automatically.
  - Modify the new connection so that it also uses the address **10.0.X.1/24**.
  - Configure the **hosts** file so that **10.0.X.1** can be referenced as "myhost".
  - Set the host name to **serverX.example.com**.

Parameter

Setting



| Parameter       | Setting        |
|-----------------|----------------|
| Connection name | review         |
| IP address      | 172.25.X.11/16 |
| Gateway address | 172.25.0.1     |
| DNS address     | 172.25.254.254 |

- e. Create a new static network connection using the settings in the table. Be sure to replace the *x* with the correct number for your systems.

```
[root@serverX ~]# nmcli con add con-name review
ifname eth0 type ethernet ip4 172.25.X.11/16 gw4
172.25.0.1

[root@serverX ~]# nmcli con mod "review" ipv4.dns
172.25.254.254
```

- f. Configure the new connection to be autostarted. Other connections should not start automatically.

```
[root@serverX ~]# nmcli con mod "review"
connection.autoconnect yes

[root@serverX ~]# nmcli con mod "System eth0"
connection.autoconnect no
```

- g. Modify the new connection so that it also uses the address 10.0.X.1/24.

```
[root@serverX ~]# nmcli con mod "review"
+ipv4.addresses 10.0.X.1/24
```

Or alternately:

```
[root@serverX ~]# echo "IPADDR1=10.0.X.1" >>
/etc/sysconfig/network-scripts/ifcfg-review

[root@serverX ~]# echo "PREFIX1=24" >>
/etc/sysconfig/network-scripts/ifcfg-review
```

- h. Configure the **hosts** file so that 10.0.X.1 can be referenced as "myhost".

```
[root@serverX ~]# echo "10.0.X.1 myhost" >>
/etc/hosts
```

- i. Set the host name to server*x*.example.com.

```
[root@serverX ~]# hostnamectl set-hostname
serverX.example.com
```

14. Synchronize the `/etc` directory tree on serverX to the `/configbackup` directory on serverX.

. To be able to create the target directory `/configbackup`, switch to the root user account with the `su` command.

```
[student@serverX ~]$ su -
```

```
Password: redhat
```

```
[root@desktopX ~]#
```

- a. Create the target directory for the configuration files on serverX.

```
[root@serverX ~]# mkdir /configbackup
```

- b. Use the `rsync` command to synchronize the `/etc` directory tree on serverX to the `/configsync` directory on serverX. Be aware that only the root user can read all the content in the `/etc` directory on serverX.

```
[root@serverX ~]# rsync -av /etc /configbackup
```

```
...
```

15. Create an archive named `/root/configuration-backup-server.tar.gz` with the `/configbackup` directory as content.

. Store the `/configbackup` directory in the `/root/configuration-backup-server.tar.gz` archive.

```
[root@serverX ~]# tar czf
/root/configuration-backup-server.tar.gz
/configbackup
```

16. To prepare the archived directory tree for comparison with the currently actively used configuration files on serverX, extract the contents of the `/root/configuration-backup-server.tar.gz` archive to the `/tmp/configcompare/` directory on serverX.

Show Solution

17. Perform the following tasks on the serverX machine:

- Use `ssh` to execute the `hostname` command on desktopX as user student. Send the output of the `hostname` command to the `/tmp/scpfile.txt` file on desktopX.

- o Use **scp** to copy the `/tmp/scpfile.txt` file from desktopX to `/home/student/scpfile.txt`.
- b. Use **ssh** to execute the **hostname** command on desktopX as user student. Send the output of the **hostname** command to the `/tmp/scpfile.txt` file on desktopX.

```
[root@serverX ~]# ssh student@desktopX
'hostname >/tmp/scpfile.txt'
```

- c. Use **scp** to copy the `/tmp/scpfile.txt` file from desktopX to `/home/student/scpfile.txt` on the serverX machine.

```
[root@serverX ~]# scp
root@desktopX:/tmp/scpfile.txt /home/student/
```

18. Create the file `/etc/yum.repos.d/localupdates.repo` to enable the “Updates” repository found on the content machine. It should access content found at the following URL: `http://content.example.com/rhel7.0/x86_64/errata`. Do *not* check GPG signatures.

Create the file `/etc/yum.repos.d/localupdates.repo` with the following content:

```
[updates]
name=Red Hat Updates
baseurl=http://content.example.com/rhel7.0/x86_64/errata
enabled=1
gpgcheck=0
```

19. Configure serverX to adhere to specific software requirements:

- o The **kernel** package should be updated to the latest version.
- o The **xsane-gimp** package should be installed.
- o The **rht-system** package should be installed.
- o For security reasons, serverX should not have the **wvdial** package installed.

- d. Update the **kernel** package.

```
yum update kernel
```

- e. Install the **xsane-gimp** package.

```
yum install xsane-gimp
```

- f. Install the **rht-system** package.

```
yum install rht-system
```

- g. For security reasons, serverX should not have the **wvdial** package installed.

```
yum remove wvdial
```

20. Generate a disk usage report with the **du** command of the **/usr/share/fonts** directory on serverX and save the result in the **/home/student/dureport.txt** file.

```
[root@serverX ~]# du
/usr/share/fonts >/home/student/dureport.txt
```

21. Identify and mount a newly added file system by UUID on the **/mnt/datadump** directory on serverX.

. Identify the newly added file system with the **blkid** command on serverX.

```
[root@serverX ~]# blkid
```

```
/dev/xvda1: UUID="46f543fd-78c9-4526-a857-244811be2d88" TYPE="xfs"
```

```
/dev/xvdj1: UUID="a84f6842-ec1d-4f6d-b767-b9570f9fcdc0" TYPE="xfs"
```

- a. Create the mount point **/mnt/datadump** on serverX.

```
[root@serverX ~]# mkdir /mnt/datadump
```

- b. Mount the file system by UUID on the **/mnt/datadump** directory of the serverX machine.

```
[root@serverX ~]# mount
UUID="a84f6842-ec1d-4f6d-b767-b9570f9fcdc0"
/mnt/datadump
```

22. Create the soft link **/root/mydataspace**, which points to the **/mnt/datadump** directory on serverX.

```
[root@serverX ~]# ln -s /mnt/datadump
/root/mydataspace
```

23. Record the command to find all soft links on serverX that have **data** as part of their name in the **/home/student/find.txt** file.

```
[root@serverX ~]# echo "find / -type l -name
'*data*' " >/home/student/find.txt
```