



redhat.®

RED HAT TRAINING

Red Hat System Administration III Student Workbook

RH254-RHEL7-en-1-20140711
MAN-RH254SKE-R2

**RED HAT®
TRAINING**



Comprehensive, hands-on training that solves real world problems

Red Hat System Administration III

Student Workbook

**RED HAT SYSTEM
ADMINISTRATION
III**

Red Hat Enterprise Linux 7 RH254

Red Hat System Administration III

Edition 1

Authors: Wander Boessenkool, Chen Chang, Will Dinyes, George Hacker,
Rudolf Kastl, Scott McBrien, Douglas Silva, Philip Sweany
Editor: Steven Bonneville

Copyright © 2014 Red Hat, Inc.

The contents of this course and all its modules and related materials, including handouts to audience members, are Copyright © 2014 Red Hat, Inc.

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Contributors and Reviewers: Chris Negus, Bruce Wolfe, Rob Locke, Bowe Strickland, Ellen Freeman, Connie Petlitzer, Brandon Nolta, Niraj Nigam, Dirk Davidis, Forrest Taylor, Jim Rigsbee, Ricardo Taniguchi

Document Conventions	xi
Notes and Warnings	xi
Introduction	xiii
Red Hat System Administration III	xiii
Orientation to the Classroom Environment	xiv
Internationalization	xv
1. Controlling Services and Daemons	1
Controlling Services with systemctl	2
Practice: Using systemctl to Manage Services	8
Controlling the Boot Process	11
Practice: Selecting a Boot Target	15
Lab: Controlling Services and Daemons	17
2. Managing IPv6 Networking	23
Review of IPv4 Networking Configuration	24
Practice: Configuring IPv4 Networking	31
IPv6 Networking Concepts	35
Practice: Interpreting IPv6 Addresses	41
IPv6 Networking Configuration	43
Practice: Configuring IPv6 Networking	49
Lab: Managing IPv6 Networking	54
3. Configuring Link Aggregation and Bridging	59
Configuring Network Teaming	60
Practice: Configuring Network Teaming	63
Managing Network Teaming	67
Practice: Managing Network Teaming	71
Configuring Software Bridges	78
Practice: Configuring Software Bridges	80
Lab: Configuring Link Aggregation and Bridging	83
4. Network Port Security	89
Managing FirewallD	90
Practice: Configuring a Firewall	94
Managing Rich Rules	97
Practice: Writing Custom Rules	101
Masquerading and Port Forwarding	103
Practice: Forwarding a Port	106
Managing SELinux Port Labeling	107
Practice: Managing SELinux Port Labeling	110
Lab: Network Port Security	113
5. Managing DNS for Servers	119
DNS Concepts	120
Practice: DNS Resource Record	128
Configuring a Caching Nameserver	131
Practice: Configuring unbound as a Caching Nameserver	135
DNS Troubleshooting	139
Practice: Troubleshooting DNS	144
Lab: Managing DNS for Servers	147
6. Configuring Email Transmission	155
Configuring Send-only Email Service	156

Practice: Configuring Send-only Email Service	163
Lab: Configuring Email Transmission	166
7. Providing Remote Block Storage	173
iSCSI Concepts	174
Practice: iSCSI Concepts	178
Providing iSCSI Targets	180
Practice: Providing iSCSI Targets	186
Accessing iSCSI Storage	189
Practice: Accessing iSCSI Storage	191
Lab: Providing Block-based Storage	195
8. Providing File-based Storage	201
Exporting NFS File Systems	202
Practice: Exporting NFS File Systems	206
Protecting NFS Exports	209
Practice: Protecting NFS Exports	214
Providing SMB File Shares	219
Practice: Providing SMB File Shares	226
Performing a Multiuser SMB Mount	231
Practice: Performing a Multiuser SMB Mount	233
Lab: Providing File-based Storage	236
9. Configuring MariaDB Databases	247
Installing MariaDB	248
Practice: Installing MariaDB	254
Working with MariaDB Databases	258
Practice: MariaDB Commands	264
Managing Database Users and Access Rights	267
Practice: Managing Users	272
Creating and Restoring MariaDB Backups	275
Practice: Restoring a MariaDB Database from Backup	280
Lab: Configuring MariaDB Databases	282
10. Providing Apache HTTPD Web Service	289
Configuring Apache HTTPD	290
Practice: Configuring a Web Server	296
Configuring and Troubleshooting Virtual Hosts	298
Practice: Configuring a Virtual Host	301
Configuring HTTPS	304
Practice: Configuring a TLS-enabled Virtual Host	310
Integrating Dynamic Web Content	315
Practice: Configuring a Web Application	318
Lab: Providing Apache HTTPD Web Service	321
11. Writing Bash Scripts	327
Bash Shell Scripting Basics	328
Practice: Writing Bash Scripts	341
Lab: Writing Bash Scripts	345
12. Bash Conditionals and Control Structures	351
Enhancing Bash Shell Scripts with Conditionals and Control Structures	352
Practice: Enhancing Bash Shell Scripts with Conditionals and Control Structures	362
Lab: Bash Conditionals and Control Structures	368

13. Configuring the Shell Environment	375
Changing the Shell Environment	376
Practice: Working with Login and Non-Login Shells	380
Lab: Configuring the Shell Environment	382
14. Linux Containers and Docker	387
Introduction to Linux Containers	388
Using Docker	394
15. Comprehensive Review	401
Red Hat System Administration III Comprehensive Review	402
Lab: Comprehensive Review of System Administration III	405

Document Conventions

Notes and Warnings



Note

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



References

"References" describe where to find external documentation relevant to a subject.



Important

"Important" boxes detail things that are easily missed: configuration changes that only apply to the current session, or services that need restarting before an update will apply. Ignoring a box labeled "Important" will not cause data loss, but may cause irritation and frustration.



Warning

"Warnings" should not be ignored. Ignoring warnings will most likely cause data loss.

Introduction

Red Hat System Administration III

Red Hat System Administration III (RH254) is designed for experienced Linux system administrators with the Red Hat Certified System Administrator (RHCSA) certification or equivalent skills. This course is intended to help students broaden their ability to administer Linux systems at an enterprise level.

Red Hat System Administration III is focused on deployment and management of network servers running caching DNS service, MariaDB, Apache HTTPD, Postfix SMTP nullclients, network file sharing with NFS and SMB, iSCSI initiators and targets, advanced networking and firewall configurations, and the use of Bash shell scripting to help automate, configure, and troubleshoot the system. Through the lectures and hands-on labs in this course, students who have already earned the RHCSA certification will be exposed to all competencies covered by the RHCE exam. This course does not cover all RHCSA tasks. Students who need a review of RHCSA skills are encouraged to attend *RHCSA Rapid Track* (RH199).

Course Objectives

- Build skills needed to be a successful senior Red Hat Enterprise Linux system administrator responsible for managing network services
- Prepare for attendance in an RHCE certification exam

Audience

- Experienced Linux administrators wishing to learn automation techniques through scripting, deployment of key network services and how to manage features of Red Hat Enterprise Linux.

Prerequisites

- RHCSA certification or equivalent experience.

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 hostnames `desktopX.example.com` and `serverX.example.com`, where the *X* in the computers' hostnames 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*.

In a live instructor-led classroom, students will be assigned a physical computer ("foundationX") which will be used to access these two machines. The **desktop** and **server** systems are virtual machines running on that host. Students should log in to this machine as user *kiosk* with the password *redhat*.

On foundationX, a special command called **rht-vmctl** is used to work with the **desktop** and **server** machines. The commands in the table below should be run as the *kiosk* user on foundationX, and can be used with **server** (as in the examples) or **desktop**.

rht-vmctl commands

Action	Command
Start server machine	rht-vmctl start server
View "physical console" to log in and work with server machine	rht-vmctl view server
Reset server machine to its previous state and restarts virtual machine	rht-vmctl reset server

At the start of a lab exercise, if the instruction "reset your server" appears, that means the command **rht-vmctl reset server** should be run in a prompt on the foundationX system as user *kiosk*. Likewise, if the instruction "reset your desktop" appears, that means the command **rht-vmctl reset desktop** should be run on foundationX as user *kiosk*.

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 which acts as a router for the classroom networks and which provides DNS, DHCP, HTTP, and other content services, `classroom.example.com`.

Classroom Machines

Machine name	IP addresses	Role
<code>desktopX.example.com</code>	<code>172.25.X.10</code>	Student "client" computer
<code>serverX.example.com</code>	<code>172.25.X.11</code>	Student "server" computer
<code>classroom.example.com</code>	<code>172.25.254.254</code>	Classroom utility server

Internationalization

Language support

Red Hat Enterprise Linux 7 officially supports twenty-two 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 on 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.

These settings affect the GNOME desktop environment and any applications, 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 login 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:

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

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

Input method settings

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.

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 on 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. Firstly, it indicates which input method is active. Secondly, it acts as a menu which 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 **DownArrow** and **UpArrow** 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

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.

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-locales LANG=fr_FR.utf8
```

In GNOME, an administrative user can change this setting from **Region & Language** and clicking on 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.

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`.

Language Codes Reference

Language Codes

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

CONTROLLING SERVICES AND DAEMONS

Overview	
Goal	To review how to manage services and the boot-up process using <code>systemctl</code> .
Objectives	<ul style="list-style-type: none">• Manage how <code>systemd</code> starts up system daemons and network services, using <code>systemctl</code>.• Control and troubleshoot system boot using <code>systemd</code> targets.
Sections	<ul style="list-style-type: none">• Controlling Services with <code>systemctl</code> (and Practice)• Controlling the Boot Process (and Practice)
Lab	<ul style="list-style-type: none">• Controlling Services and Daemons

Controlling Services with `systemctl`

Objectives

After completing this section, students should be able to:

- List system daemons and network services started by the **systemd** service and socket units.
- Control system daemons and network services using **systemctl**.

Introduction to `systemd`

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.

Daemons are processes that wait or run in the background performing various tasks. To listen for connections, a daemon uses a *socket*. Sockets may be created by daemons or may be separated from the daemon and be created by another process, such as `systemd`, which then passes the socket to the daemon when a connection is established by a 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 (for example, to configure network interfaces), which does not involve leaving a daemon process running afterward.

A bit of history

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. Frequently used daemons were started on systems at boot time with *System V* and Linux Standard Base (*LSB*) init scripts. Less frequently used daemons were started on demand by another service, such as **initd** or **xinetd**. 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.
- On-demand starting of daemons without requiring a separate service.
- Automatic service dependency management prevents long timeouts, such as not starting a network service when the network is not available.
- A method of tracking related processes together using Linux control groups.



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**

The **systemctl** may abbreviate or "ellipsize" unit names, process tree entries, and unit descriptions unless run with the **-l** option.

Some common unit types are listed as follows:

- *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.
- *Socket units* have a `.socket` extension and represent interprocess communication (IPC) sockets. 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.
- *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.

Service states

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 (running)	Running with one or more continuing processes.
active (exited)	Successfully completed a one-time configuration.
active (waiting)	Running but waiting for an event.
inactive	Not running.

Keyword:	Description:
enabled	Will be started at boot time.
disabled	Will not be started at boot time.
static	Cannot 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

- Query the state of all units to verify a system startup.

```
[root@serverX ~]# systemctl
```

- Query the state of only the service units.

```
[root@serverX ~]# systemctl --type=service
```

- Investigate any units which are in a failed or maintenance state. Optionally, add the **-l** option to show the full output.

```
[root@serverX ~]# systemctl status rngd.service -l
```

- 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:

```
[root@serverX ~]# systemctl is-active sshd
[root@serverX ~]# systemctl is-enabled sshd
```

- List the active state of all loaded units. Optionally, limit the type of unit. The **--all** option will add inactive units.

```
[root@serverX ~]# systemctl list-units --type=service
[root@serverX ~]# systemctl list-units --type=service --all
```

- View the enabled and disabled settings for all units. Optionally, limit the type of unit.

```
[root@serverX ~]# systemctl list-unit-files --type=service
```

- View only failed services.

```
[root@serverX ~]# systemctl --failed --type=service
```

Starting and stopping system daemons on a running system

Starting, stopping, restarting, reloading, and verifying status are common actions performed when administering services.

- View the status of the **sshd** service.

```
[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
```

- Verify that the process is running.

```
[root@serverX ~]# ps -up 1073
USER      PID %CPU %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
root       1073  0.1  0.0  82992  3612 ?        Ss   15:15   0:00 /usr/sbin/sshd -D
```

- Stop the service and verify the status.

```
[root@serverX ~]# systemctl stop sshd.service
[root@serverX ~]# systemctl status sshd.service
sshd.service - OpenSSH server daemon
  Loaded: loaded (/usr/lib/systemd/system/sshd.service; enabled)
  Active: inactive (dead) since Thu 2014-02-27 18:51:39 EST; 2s ago
  Main PID: 1073 (code=exited, status=0/SUCCESS)
```

- Start the service and view the status. The process ID has changed.

```
[root@serverX ~]# systemctl start sshd.service
[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 18:52:39 EST; 2s ago
  Main PID: 1253 (sshd)
  CGroup: /system.slice/sshd.service
          └─1253 /usr/sbin/sshd -D
```

- Stop, then start, the service in a single command.

```
[root@serverX ~]# systemctl restart sshd.service
[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 18:54:39 EST; 2s ago
  Main PID: 1268 (sshd)
  CGroup: /system.slice/sshd.service
          └─1268 /usr/sbin/sshd -D
```

- Issue instructions for a service to read and reload its configuration file without a complete stop and start. The process ID will not change.

```
[root@serverX ~]# systemctl reload sshd.service
[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 18:55:09 EST; 32s ago
     Main PID: 1268 (sshd)
        CGroup: /system.slice/sshd.service
                └─1268 /usr/sbin/sshd -D
```

Unit dependencies

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.

The **systemctl list-dependencies UNIT** command can be used to display a tree of other units which must be started in conjunction with a specific unit. The **--reverse** option to this command will show what units need to have the specified unit started in order to run.

Masking services

A system may have conflicting services installed for a certain function, such as firewalls (iptables and firewalld). To prevent an administrator from accidentally starting a service, a service may be *masked*. Masking creates a link in the configuration directories so that if the service is started, nothing will happen.

```
[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'
```



Important

A disabled service will not be started automatically at boot or by other unit files, but can be started manually. A masked service cannot be started manually or automatically.

Enabling system daemons to start or stop at boot

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.

- View the status of a service.

```
[root@serverX ~]# systemctl status sshd.service
```

- Disable the service and verify the status. Note that disabling a service does not stop the service.

```
[root@serverX ~]# systemctl disable sshd.service
[root@serverX ~]# systemctl status sshd.service
```

- Enable the service and verify the status.

```
[root@serverX ~]# systemctl enable sshd.service
[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.

Command:	Task:
systemctl status UNIT	View detailed information about a unit state.
systemctl stop UNIT	Stop a service on a running system.
systemctl start UNIT	Start a service on a running system.
systemctl restart UNIT	Restart a service on a running system.
systemctl reload UNIT	Reload configuration file of a running service.
systemctl mask UNIT	Completely disable a service from being started, both manually and at boot.
systemctl unmask UNIT	Make a masked service available.
systemctl enable UNIT	Configure a service to start at boot time.
systemctl disable UNIT	Disable a service from starting at boot time.
systemctl list-dependencies UNIT	List units which are required and wanted by the specified unit.

R

References

systemd(2), **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/>

Practice: Using `systemctl` to Manage Services

Guided exercise

In this lab, you will manage a service unit that is already installed on the system.

Resources	
Machines:	serverX

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 the **systemctl restart** and **systemctl reload** commands.
 - 1.1. Display the status of the **sshd** service. Note the process ID of the daemon.

```
[student@serverX ~]$ sudo systemctl status sshd
shd.service - OpenSSH server daemon
Loaded: loaded (/usr/lib/systemd/system/ssh.service; enabled)
Active: active (running) since Tue 2014-06-10 09:09:50 EDT; 35min ago
Process: 1061 ExecStartPre=/usr/sbin/ssh-keygen (code=exited, status=0/SUCCESS)
Main PID: 1077 (sshd)
CGroup: /system.slice/ssh.service
└─1077 /usr/sbin/sshd -D
...
```

- 1.2. Restart the **sshd** service and view the status. The process ID of the daemon has changed.

```
[student@serverX ~]$ sudo systemctl restart sshd
[student@serverX ~]$ sudo systemctl status sshd
ssh.service - OpenSSH server daemon
Loaded: loaded (/usr/lib/systemd/system/ssh.service; enabled)
Active: active (running) since Tue 2014-06-10 09:48:00 EDT; 1s ago
Process: 2851 ExecStartPre=/usr/sbin/ssh-keygen (code=exited, status=0/SUCCESS)
Main PID: 2852 (sshd)
CGroup: /system.slice/ssh.service
└─2852 /usr/sbin/sshd -D
...
```

- 1.3. Reload the **sshd** service and view the status. The process ID of the daemon has not changed and connections have not been interrupted.

```
[student@serverX ~]$ sudo systemctl reload sshd
[student@serverX ~]$ sudo systemctl status sshd
sshd.service - OpenSSH server daemon
Loaded: loaded (/usr/lib/systemd/system/sshd.service; enabled)
Active: active (running) since Tue 2014-06-10 09:48:00 EDT; 12s ago
Process: 2860 ExecReload=/bin/kill -HUP $MAINPID (code=exited, status=0/SUCCESS)
Process: 2851 ExecStartPre=/usr/sbin/sshd-keygen (code=exited, status=0/SUCCESS)
Main PID: 2852 (sshd)
CGroup: /system.slice/sshd.service
└─2852 /usr/sbin/sshd -D
...
```

- 2. Verify that the **chronyd** service is running.

```
[student@serverX ~]$ sudo systemctl status chronyd
chronyd.service - NTP client/server
Loaded: loaded (/usr/lib/systemd/system/chronyd.service; enabled)
Active: active (running) since Tue 2014-06-10 09:09:44 EDT; 43min ago
Process: 490 ExecStartPost=/usr/libexec/chrony-helper add-dhclient-servers (code=exited, status=0/SUCCESS)
Process: 450 ExecStart=/usr/sbin/chronyd -u chrony $OPTIONS (code=exited, status=0/SUCCESS)
Main PID: 467 (chronyd)
CGroup: /system.slice/chronyd.service
└─467 /usr/sbin/chronyd -u chrony
...
```

- 3. Stop the **chronyd** service and view the status.

```
[student@serverX ~]$ sudo systemctl stop chronyd
[student@serverX ~]$ sudo systemctl status chronyd
chronyd.service - NTP client/server
Loaded: loaded (/usr/lib/systemd/system/chronyd.service; enabled)
Active: inactive (dead) since Tue 2014-06-10 09:53:36 EDT; 4s ago
Process: 490 ExecStartPost=/usr/libexec/chrony-helper add-dhclient-servers (code=exited, status=0/SUCCESS)
Process: 450 ExecStart=/usr/sbin/chronyd -u chrony $OPTIONS (code=exited, status=0/SUCCESS)
Main PID: 467 (code=exited, status=0/SUCCESS)
...
```

- 4. Determine if the **chronyd** service is enabled to start at system boot.

```
[student@serverX ~]$ sudo systemctl is-enabled chronyd
enabled
```

- 5. Reboot the system.

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

- 6. Log in to the serverX system and view the status of the **chronyd** service.

```
[student@serverX ~]$ sudo systemctl status chronyd
chronyd.service - NTP client/server
Loaded: loaded (/usr/lib/systemd/system/chronyd.service; enabled)
Active: active (running) since Tue 2014-06-10 09:55:22 EDT; 9min ago
Process: 487 ExecStartPost=/usr/libexec/chrony-helper add-dhclient-servers
        (code=exited, status=0/SUCCESS)
Process: 451 ExecStart=/usr/sbin/chronyd -u chrony $OPTIONS (code=exited,
        status=0/SUCCESS)
Main PID: 474 (chronyd)
CGroup: /system.slice/chronyd.service
└─474 /usr/sbin/chronyd -u chrony
...
```

- ❑ 7. Disable the **chronyd** service so that it does not start at system boot, then view the status of the service.

```
[student@serverX ~]$ sudo systemctl disable chronyd
[student@serverX ~]$ sudo systemctl status chronyd
chronyd.service - NTP client/server
Loaded: loaded (/usr/lib/systemd/system/chronyd.service; disabled)
Active: active (running) since Tue 2014-06-10 09:55:22 EDT; 9min ago
Main PID: 474 (chronyd)
CGroup: /system.slice/chronyd.service
└─474 /usr/sbin/chronyd -u chrony
...
```

- ❑ 8. Reboot the system.

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

- ❑ 9. Log in to the serverX system again and view the status of the **chronyd** service.

```
[student@serverX ~]$ sudo systemctl status chronyd
chronyd.service - NTP client/server
Loaded: loaded (/usr/lib/systemd/system/chronyd.service; disabled)
Active: inactive (dead)
```


Controlling the Boot Process

Objectives

After completing this section, students should be able to influence the boot process and be able to repair common boot issues using **systemd** targets.

Selecting a systemd target

A **systemd** target is a set of **systemd** units that should be started to reach a desired state. Important targets are listed in the following table.

Target	Purpose
graphical.target	System supports multiple users, graphical and text-based logins.
multi-user.target	System supports multiple users, text-based logins only.
rescue.target	sulogin prompt, basic system initialization completed.
emergency.target	sulogin prompt, initramfs pivot complete and system root mounted on / read-only.

It is possible for a target to be a part of another target; for example, the **graphical.target** includes **multi-user.target**, which in turn depends on **basic.target** and others. These dependencies can be viewed from the command line with the following command:

```
[root@serverX ~]# systemctl list-dependencies graphical.target | grep target
```

An overview of all available targets can be viewed with:

```
[root@serverX ~]# systemctl list-units --type=target --all
```

An overview of all targets installed on disk can be viewed with:

```
[root@serverX ~]# systemctl list-unit-files --type=target --all
```

Selecting a target at runtime

On a running system, administrators can choose to switch to a different target using the **systemctl isolate** command; for example, **systemctl isolate multi-user.target**.



Note

Not all targets can be isolated. Only targets that have **AllowIsolate=yes** set in their unit files can be isolated; for example, the **graphical.target** target can be isolated, but the **cryptsetup.target** target cannot.

Setting a default target

When the system starts, and control is passed over to **systemd** from the **initramfs**, **systemd** will try to activate the **default.target** target. Normally the **default.target** target will

be a symbolic link (in `/etc/systemd/system/`) to either `graphical.target` or `multi-user.target`.

The `systemctl` tool provides two commands to manage the link: `get-default` and `set-default`.

```
[root@serverX ~]# systemctl get-default
multi-user.target
[root@serverX ~]# systemctl set-default graphical.target
rm '/etc/systemd/system/default.target'
ln -s '/usr/lib/systemd/system/graphical.target' '/etc/systemd/system/default.target'
[root@serverX ~]# systemctl get-default
graphical.target
```

Selecting a different target at boot time

To select a different target at boot time, a special option can be appended to the kernel command line from the boot loader: `systemd.unit=`. For example, to boot the system into a rescue shell, pass the following option at the interactive boot loader menu:

```
systemd.unit=rescue.target
```

To use this method of selecting a different target, use the following procedure for Red Hat Enterprise Linux 7 systems:

1. (Re)boot the system.
2. Interrupt the boot loader menu countdown by pressing any key.
3. Move the cursor to the entry to be started.
4. Press `e` to edit the current entry.
5. Move the cursor to the line that starts with `linux16`. This is the kernel command line.
6. Append `systemd.unit=desired.target`.
7. Press `Ctrl+x` to boot with these changes.

Recovering the root password

Recovering the `root` password is a trivial task while still logged in as an administrator or a user with full `sudo` access, but is slightly more involved when an administrator is not logged in. In the latter situation, the administrator could boot from a Live CD, mount the root file system from there, and edit `/etc/shadow`. Administrators should also be able to perform root password recovery without the use of external media.



Note

On Red Hat Enterprise Linux 6 and earlier, an administrator could boot the system into `runlevel 1`, and be presented with a root prompt. The closest analogs to `runlevel 1` on a Red Hat Enterprise Linux 7 machine are the `rescue.target` and `emergency.target` targets, both of which require the `root` password to log in.

On Red Hat Enterprise Linux 7, it is possible to have the scripts that run from the **initramfs** pause at certain points, provide a **root** shell, and then continue when that shell exits. While this is mostly meant for debugging, it can also be used to recover a lost **root** password:

1. Reboot the system.
2. Interrupt the boot loader countdown by pressing any key.
3. Move the cursor to the entry that needs to be booted.
4. Press **e** to edit the selected entry.
5. Move the cursor to the kernel command line (the line that starts with **linux16**).
6. Append **rd.break** (this will break just before control is handed from the **initramfs** to the actual system).



Note

The **initramfs** prompt will show up on whatever console is specified *last* on the kernel command line.

7. Press **Ctrl+x** to boot with the changes.

At this point, a **root** shell will be presented, with the root file system for the actual system mounted read-only on **/sysroot**.



Important

SELinux is not yet enabled at this point, so any new files being created will not have an SELinux context assigned to them. Keep in mind that some tools (such as **passwd**) first create a new file, then move it in place of the file they are intended to edit, effectively creating a new file without an SELinux context.

To recover the **root** password from this point, use the following procedure:

1. Remount **/sysroot** as read-write.

```
switch_root:/# mount -o remount,rw /sysroot
```

2. Switch into a **chroot** jail, where **/sysroot** is treated as the root of the file system tree.

```
switch_root:/# chroot /sysroot
```

3. Set a new root password:

```
sh-4.2# passwd root
```

4. Make sure that all unlabeled files (including **/etc/shadow** at this point) get relabeled during boot.

```
sh-4.2# touch /.autorelabel
```

5. Type **exit** twice. The first will exit the **chroot** jail, and the second will exit the **initramfs** debug shell.

At this point, the system will continue booting, perform a full SELinux relabel, then reboot again.

Diagnose and repair systemd boot issues

If there are problems during the starting of services, there are a few tools available to system administrators that can help with debugging and/or troubleshooting:

Early debug shell

By running **systemctl enable debug-shell.service**, a **root** shell will be spawned on **TTY9 (Ctrl+Alt+F9)** early during the boot sequence. This shell is automatically logged in as root so that an administrator can use some of the other debugging tools while the system is still booting.



Warning

Do not forget to disable the **debug-shell.service** service when debugging is complete, as it leaves an unauthenticated root shell open to anyone with local console access.

Emergency and rescue targets

By appending either **systemd.unit=rescue.target** or **systemd.unit=emergency.target** to the kernel command line from the boot loader, the system will spawn into a special rescue or emergency shell instead of starting normally. Both of these shells require the **root** password. The **emergency** target keeps the root file system mounted read-only, while **rescue.target** waits for **sysinit.target** to complete first so that more of the system will be initialized, for example, logging, file systems, etc. Exiting from these shells will continue with the regular boot process.

Stuck jobs

During startup, **systemd** spawns a number of jobs. If some of these jobs cannot complete, they will block other jobs from running. To inspect the current job list, an administrator can use the command **systemctl list-jobs**. Any jobs listed as **running** must complete before the jobs listed as **waiting** can continue.

R

References

systemd.target(5), **systemd.special(7)**, **sulogin(8)**, **sushell(8)**, and **systemctl(1)** man pages

/usr/lib/systemd/system/debug-shell.service

Practice: Selecting a Boot Target

Guided exercise

In this lab, you will configure your **serverX** system to boot into different targets.

Resources:

Machines:	serverX
------------------	----------------

Outcome:

A system booted into different targets.

Before you begin...

- Reset your **serverX** system.

- 1. On your **serverX** system, switch to the **multi-user** target manually without rebooting.

- 1.1.

```
[student@serverX ~]$ sudo systemctl isolate multi-user.target
```

- 2. Log into a text-based console as **root**.

- 3. Configure your **serverX** to automatically boot into the **multi-user** target after a reboot, then reboot your **serverX** system to verify.

- 3.1.

```
[root@serverX ~]# systemctl set-default multi-user.target
rm '/etc/systemd/system/default.target'
ln -s '/usr/lib/systemd/system/multi-user.target' '/etc/systemd/system/default.target'
```

- 3.2.

```
[root@serverX ~]# systemctl reboot
```

- 4. Reboot your **serverX** system, then from within the boot loader menu, boot into the **rescue** target.

- 4.1. Reboot your **serverX** machine.

- 4.1.

```
[root@serverX ~]# systemctl reboot
```

- 4.2. Interrupt the boot loader when the menu appears by pressing any key.

- 4.3. Move the selection to the default entry (the first one) using the cursor keys.

- 4.4. Press **e** to edit the current entry.

- 4.5. Move the cursor to the line that starts with **linux16**.

- 4.6. Move the cursor to the end of the line (using the **End** key), and append the following text:

```
systemd.unit=rescue.target
```

- 4.7. Press **Ctrl+x** to boot using the modified configuration.
- 4.8. When prompted for the **root** password, enter **redhat**.
- 5. Set the default **systemd** target back to the graphical target.

```
[root@serverX ~]# systemctl set-default graphical.target
```

- 6. Press **Ctrl+d** to continue booting into the (new) default target.

Lab: Controlling Services and Daemons

Performance checklist

In this lab, you will change the default target and enable and verify that a service starts upon boot.

Resources:	
Machines:	serverX

Outcomes:

Students will configure the serverX.example.com machine to boot to a state supporting multiple users with both graphical and text-based logins. Students will also configure the rsyslog service to start at boot time.

Before you begin...

- Reset your **serverX** system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab systemd setup
```

You have received a request from users to enable graphical login on **serverX**. The users have also asked you to look into the logging service on the system, since the log files do not seem to be getting populated. You will make the necessary configuration changes and then reboot the system.

After the reboot, verify the system state and rsyslog service status to make sure everything is working as expected. After completing your work, run the command **lab systemd grade** on **serverX** to verify the results.

1. Determine if the system currently supports graphical login on boot. If not, configure it to support graphical and text-based logins upon boot.
2. Determine if the rsyslog process is running. If not, verify if it is configured to start on boot, and fix it if it is not.
3. Start the rsyslog service manually to make sure that it will start properly.
4. Reboot the system.
5. Log in to the serverX system again and switch to user **root**. Verify that the system boots to the desired system state and that the rsyslog service is running properly.
6. Verify your work by running **lab systemd grade** on **serverX**.

Solution

In this lab, you will change the default target and enable and verify that a service starts upon boot.

Resources:

Machines:

serverX

Outcomes:

Students will configure the serverX.example.com machine to boot to a state supporting multiple users with both graphical and text-based logins. Students will also configure the rsyslog service to start at boot time.

Before you begin...

- Reset your **serverX** system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab systemd setup
```

You have received a request from users to enable graphical login on **serverX**. The users have also asked you to look into the logging service on the system, since the log files do not seem to be getting populated. You will make the necessary configuration changes and then reboot the system.

After the reboot, verify the system state and rsyslog service status to make sure everything is working as expected. After completing your work, run the command **lab systemd grade** on **serverX** to verify the results.

1. Determine if the system currently supports graphical login on boot. If not, configure it to support graphical and text-based logins upon boot.
 - 1.1. Switch to the root user on the serverX system.

```
[student@serverX ~]$ sudo -i
```

- 1.2. Determine the state that the system is currently configured to boot to.

```
[root@serverX ~]# systemctl get-default
multi-user.target
```

- 1.3. Set the system state to support graphical login.

```
[root@serverX ~]# systemctl set-default graphical.target
rm '/etc/systemd/system/default.target'
ln -s '/usr/lib/systemd/system/graphical.target' '/etc/systemd/system/
default.target'
```

2. Determine if the rsyslog process is running. If not, verify if it is configured to start on boot, and fix it if it is not.
 - 2.1. Determine if the rsyslog process is running.

```
[root@serverX ~]# ps -lef | grep [r]syslog
[root@serverX ~]#
```

- 2.2. Check if the rsyslog service is enabled.

```
[root@serverX ~]# systemctl is-enabled rsyslog.service
disabled
```

- 2.3. Enable the rsyslog service.

```
[root@serverX ~]# systemctl enable rsyslog.service
ln -s '/usr/lib/systemd/system/rsyslog.service' '/etc/systemd/system/multi-user.target.wants/rsyslog.service'
```

3. Start the rsyslog service manually to make sure that it will start properly.

- 3.1. Start the rsyslog service.

```
[root@serverX ~]# systemctl start rsyslog.service
[root@serverX ~]#
```

- 3.2. Verify the rsyslog service started without issues.

```
[root@serverX ~]# systemctl status rsyslog.service
rsyslog.service - System Logging Service
   Loaded: loaded (/usr/lib/systemd/system/rsyslog.service; enabled)
   Active: active (running) since Tue 2014-05-27 00:55:57 EDT; 16s ago
     Main PID: 1910 (rsyslogd)
        CGroup: /system.slice/rsyslog.service
                └─1910 /usr/sbin/rsyslogd -n

May 27 00:55:57 server1.example.com systemd[1]: Started System Logging Service.
```

4. Reboot the system.

```
[root@serverX ~]# systemctl reboot
```

5. Log in to the serverX system again and switch to user **root**. Verify that the system boots to the desired system state and that the rsyslog service is running properly.

- 5.1. Verify the system state. The graphical target should be in state active on a fully booted up system.

```
[root@serverX ~]# systemctl status graphical.target
graphical.target - Graphical Interface
   Loaded: loaded (/lib/systemd/system/graphical.target; enabled)
   Active: active since Tue 2014-05-27 01:11:26 EDT; 1min 52s ago
     Docs: man:systemd.special(7)
   ...
```

- 5.2. Verify that the rsyslog service is running properly.

```
[root@serverX ~]# systemctl status rsyslog.service
rsyslog.service - System Logging Service
   Loaded: loaded (/usr/lib/systemd/system/rsyslog.service; enabled)
   Active: active (running) since Tue 2014-05-27 01:11:22 EDT; 1min 48s ago
 Main PID: 570 (rsyslogd)
    CGroup: /system.slice/rsyslog.service
           └─570 /usr/sbin/rsyslogd -n

May 27 01:11:22 localhost systemd[1]: Started System Logging Service.
```

6. Verify your work by running **lab systemd grade** on **serverX**.

```
[root@serverX ~]# lab systemd grade
```

Summary

Controlling Services with `systemctl`

In this section, students learned how to:

- Determine the status of system daemons and network services started by **systemd**.
- Start, stop, and enable services using **systemctl**.

Controlling the Boot Process

In this section, students learned how to:

- Break down the Red Hat Enterprise Linux 7 boot process into four steps:
 1. Hardware (BIOS/UEFI)
 2. Boot loader (**grub2**)
 3. **kernel** and **initramfs**
 4. **systemd**
- Control the Red Hat Enterprise Linux 7 boot process by:
 - Selecting a systemd target with **systemctl** on a running system.
 - Passing arguments on the kernel command line during boot.



CHAPTER 2

MANAGING IPV6 NETWORKING

Overview	
Goal	To configure and troubleshoot basic IPv6 networking on Red Hat Enterprise Linux systems.
Objectives	<ul style="list-style-type: none">• Review how to configure IPv4 networking in RHEL 7.• Explain the basic concepts of IPv6 networking, and read and write condensed IPv6 addresses.• Configure IPv6 networking using command-line tools and configuration files.
Sections	<ul style="list-style-type: none">• Review of IPv4 Networking Configuration (and Practice)• IPv6 Networking Concepts (and Practice)• IPv6 Networking Configuration (and Practice)
Lab	<ul style="list-style-type: none">• Managing IPv6 Networking

Review of IPv4 Networking Configuration

Objectives

After completing this section, students should be able to configure IPv4 networking using **nmcli** and configuration files in the **/etc/sysconfig/network-scripts** directory.

IPv4 networking

This section assumes that students have a basic understanding of IPv4 networking concepts. In particular, students should know something about IPv4 addresses, network prefixes (and netmasks), default gateways and basic routing, network interfaces, **/etc/hosts**, and name resolution.

NetworkManager overview

In Red Hat Enterprise Linux 7, the configuration of network interfaces is managed by a system daemon called **NetworkManager**. For NetworkManager:

- A *device* is a network interface.
- A *connection* is a collection of settings that can be configured for a device.
- Only one connection is *active* for any one device at a time. Multiple connections may exist, for use by different devices or to allow a configuration to be altered for the same device.
- Each connection has a *name* or ID that identifies it.
- The persistent configuration for a connection is stored in **/etc/sysconfig/network-scripts/ifcfg-*name***, where *name* is the name of the connection (although spaces are normally replaced with underscores in the file name). This file can be edited by hand if desired.
- The **nmcli** utility can be used to create and edit connection files from the shell prompt.

Viewing networking information

The command **nmcli dev status** will show the status of all network devices:

```
[student@demo ~]$ nmcli dev status
DEVICE  TYPE      STATE      CONNECTION
eno1    ethernet  connected  eno1
eth0    ethernet  connected  static-eth0
eno2    ethernet  disconnected --
lo      loopback  unmanaged  --
```

The command **nmcli con show** will show a list of all connections. To list only the active connections, add the **--active** option.

```
[student@demo ~]$ nmcli con show
NAME                UUID                                TYPE                DEVICE
eno2                ff9f7d69-db83-4fed-9f32-939f8b5f81cd 802-3-ethernet      --
static-eth0        72ca57a2-f780-40da-b146-99f71c431e2b 802-3-ethernet      eth0
eno1                87b53c56-1f5d-4a29-a869-8a7bdaf56dfa 802-3-ethernet      eno1
[root@demo ~]# nmcli con show --active
NAME                UUID                                TYPE                DEVICE
static-eth0        72ca57a2-f780-40da-b146-99f71c431e2b 802-3-ethernet      eth0
eno1                87b53c56-1f5d-4a29-a869-8a7bdaf56dfa 802-3-ethernet      eno1
```

The **ip addr show** command displays the current configuration of network interfaces on the system. To list only a single interface, add the interface name as the last argument:

```
[student@demo ~]$ ip addr show eth0
2: eth0: <BROADCAST,MULTICAST, 1UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
    2link/ether 52:54:00:00:00:0b brd ff:ff:ff:ff:ff:ff
    3inet 172.25.0.11/16 brd 172.25.255.255 scope global eth0
        valid_lft forever preferred_lft forever
    4inet6 fe80::5054:ff:fe00:b/64 scope link
        valid_lft forever preferred_lft forever
```

- 1** An active interface is **UP**.
- 2** The **link/ether** line specifies the hardware (MAC) address of the device.
- 3** The **inet** line shows an IPv4 address, its network prefix length, and scope.
- 4** The **inet6** line shows an IPv6 address, its network prefix length, and scope.

Adding a network connection

The **nmcli con add** command is used to add new network connections. The example **nmcli con add** commands that follow assume that the name of the network connection being added is not already in use.

The following command will add a new connection for the interface **eno2**, which will get IPv4 networking information using DHCP and will autoconnect on startup. The configuration will be saved in **/etc/sysconfig/network-scripts/ifcfg-eno2** because the **con-name** is **eno2**.

```
[root@demo ~]# nmcli con add con-name eno2 type ethernet ifname eno2
```

The next example configures the **eno2** interface statically instead, using the IPv4 address and network prefix 192.168.0.5/24 and default gateway 192.168.0.254, but still autoconnects at startup and saves its configuration into the same file. The example is line-wrapped with a shell **** escape.

```
[root@demo ~]# nmcli con add con-name eno2 type ethernet ifname eno2 \
> ip4 192.168.0.5/24 gw4 192.168.0.254
```

Controlling network connections

The **nmcli con up name** command will activate the connection *name* on the network interface it is bound to. Note that the command takes the name of a *connection*, not the name of the

network interface. Remember that **nmcli con show** can be used to list the names of all available connections.

```
[root@demo ~]# nmcli con up static-eth0
```

The **nmcli dev disconnect *device*** command will disconnect the network interface *device* and bring it down. This command can be abbreviated **nmcli dev dis *device***:

```
[root@demo ~]# nmcli dev dis eth0
```



Important

Use **nmcli dev dis *device*** to deactivate a network interface.

The command **nmcli con down *name*** is normally not the best way to deactivate a network interface. This command will bring down the connection. But by default, most wired system connections are configured with **autoconnect** enabled. This activates the connection as soon as its network interface is available. Since the connection's network interface is still available, **nmcli con down *name*** will bring the interface down, but then NetworkManager will immediately bring it up again unless the connection is entirely disconnected from the interface.

Modifying network connection settings

NetworkManager connections have two kinds of settings. There are *static* connection properties, which are configured by the administrator and stored in the configuration files in **/etc/sysconfig/network-scripts/ifcfg-***. There may also be *active* connection data, which the connection gets from a DHCP server and which are not stored persistently.

To list the current settings for a connection, run the **nmcli con show *name*** command, where *name* is the name of the connection. Settings in lowercase are static properties the administrator can change; settings in all caps are active settings in temporary use for this instance of the connection.

```
[root@demo ~]# nmcli con show static-eth0
connection.id:                static-eth0
connection.uuid:              87b53c56-1f5d-4a29-a869-8a7bdaf56dfa
connection.interface-name:    --
connection.type:              802-3-ethernet
connection.autoconnect:       yes
connection.timestamp:         1401803453
connection.read-only:         no
connection.permissions:       --
connection.zone:              --
connection.master:            --
connection.slave-type:        --
connection.secondaries:       --
connection.gateway-ping-timeout: 0
802-3-ethernet.port:          --
802-3-ethernet.speed:         0
802-3-ethernet.duplex:        --
802-3-ethernet.auto-negotiate: yes
802-3-ethernet.mac-address:   CA:9D:E9:2A:CE:F0
802-3-ethernet.cloned-mac-address: --
802-3-ethernet.mac-address-blacklist: --
802-3-ethernet.mtu:           auto
802-3-ethernet.s390-subchannels: --
802-3-ethernet.s390-nettype:  --
802-3-ethernet.s390-options:  --
ipv4.method:                  manual
ipv4.dns:                     192.168.0.254
ipv4.dns-search:              example.com
ipv4.addresses:               { ip = 192.168.0.2/24, gw = 192.168.0.254 }
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
...
```

The **nmcli con mod *name*** command can be used to change the settings for a connection. These changes will also be saved in the **/etc/sysconfig/network-scripts/ifcfg-*name*** file for the connection. The different settings that are available are documented in the **nm-settings(5)** man page.

To set the IPv4 address to 192.0.2.2/24 and default gateway to 192.0.2.254 for the connection **static-eth0**:

```
[root@demo ~]# nmcli con mod static-eth0 ipv4.addresses "192.0.2.2/24 192.0.2.254"
```



Important

If a connection that got its IPv4 information from a DHCPv4 server is being changed to get it from static configuration files only, the setting **ipv4.method** should also be changed from **auto** to **manual**. Otherwise, the connection may hang or not complete successfully when it is activated, or it may get an IPv4 address from DHCP in addition to the static address.

A number of settings may have multiple values. A specific value can be added to the list or deleted from the list for a setting by adding a **+** or **-** symbol to the start of the setting name.

To add the DNS server 192.0.2.1 to the list of nameservers to use with the connection **static-eth0**:

```
[root@demo ~]# nmcli con mod static-eth0 +ipv4.dns 192.0.2.1
```

By default, changes made with **nmcli con mod name** are automatically saved to **/etc/sysconfig/network-scripts/ifcfg-name**. That file can also be manually edited with a text editor. After doing so, run **nmcli con reload** so that NetworkManager reads the configuration changes.

For backward-compatibility reasons, the directives saved in that file have different names and syntax than the **nm-settings(5)** names. The following table maps some of the key setting names to **ifcfg-*** directives.

Comparison of nm-settings and ifcfg-* Directives

nmcli con mod	ifcfg-* file	Effect
ipv4.method manual	BOOTPROTO=none	IPv4 addresses configured statically.
ipv4.method auto	BOOTPROTO=dhcp	Will look for configuration settings from a DHCPv4 server. If static addresses are also set, will not bring those up until we have information from DHCPv4.
ipv4.addresses "192.0.2.1/24 192.0.2.254"	IPADDR0=192.0.2.1 PREFIX0=24 GATEWAY0=192.0.2.254	Sets static IPv4 address, network prefix, and default gateway. If more than one is set for the connection, then instead of 0 , the ifcfg-* directives end with 1, 2, 3 and so on.
ipv4.dns 8.8.8.8	DNS0=8.8.8.8	Modify /etc/resolv.conf to use this nameserver .
ipv4.dns-search example.com	DOMAIN=example.com	Modify /etc/resolv.conf to use this domain in the search directive.
ipv4.ignore-auto-dns true	PEERDNS=no	Ignore DNS server information from the DHCP server.
connection.autoconnect yes	ONBOOT=yes	Automatically activate this connection at boot.
connection.id eth0	NAME=eth0	The name of this connection.
connection.interface-name eth0	DEVICE=eth0	The connection is bound to the network interface with this name.

nmcli con mod	ifcfg-* file	Effect
802-3-ethernet.mac-address . . .	HWADDR= . . .	The connection is bound to the network interface with this MAC address.



Important

Because NetworkManager tends to directly modify the `/etc/resolv.conf` file, direct edits to that file may be overwritten.

To change settings in that file, it is better to set **DNSn** and **DOMAIN** directives in the relevant `/etc/sysconfig/network-scripts/ifcfg-*` files.

Deleting a network connection

The `nmcli con del name` command will delete the connection named *name* from the system, disconnecting it from the device and removing the file `/etc/sysconfig/network-scripts/ifcfg-name`.

Modifying the system host name

The `hostname` command displays or temporarily modifies the system's fully qualified host name.

```
[root@demo ~]# hostname
demo.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@demo ~]# hostnamectl set-hostname demo.example.com
[root@demo ~]# hostnamectl status
  Static hostname: demo.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:GA:server
    Kernel: Linux 3.10.0-121.el7.x86_64
  Architecture: x86_64
[root@demo ~]# cat /etc/hostname
demo.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.

Summary of commands

The following table is a list of key commands discussed in this section.

Command	Purpose
nmcli dev status	Show the NetworkManager status of all network interfaces.
nmcli con show	List all connections.
nmcli con show <i>name</i>	List the current settings for the connection <i>name</i> .
nmcli con add con-name <i>name</i> ...	Add a new connection named <i>name</i> .
nmcli con mod <i>name</i> ...	Modify the connection <i>name</i> .
nmcli con reload	Tell NetworkManager to reread the configuration files (useful after they have been edited by hand).
nmcli con up <i>name</i>	Activate the connection <i>name</i> .
nmcli dev dis <i>dev</i>	Deactivate and disconnect the current connection on the network interface <i>dev</i> .
nmcli con del <i>name</i>	Delete the connection <i>name</i> and its configuration file.
ip addr show	Show the current network interface address configuration.
hostnamectl set-hostname ...	Persistently set the host name on this system.

R

References

NetworkManager(8), **nmcli**(1), **nmcli-examples**(5), **nm-settings**(5), **hostnamectl**(1), **resolv.conf**(5), **hostname**(5), **ip**(8), and **ip-address**(8) man pages

Practice: Configuring IPv4 Networking

Guided exercise

In this lab, you will configure a network interface with a static IPv4 address. Once the interface is configured, you will confirm that it works and identify other IPv4 nodes on the local network. You will also explore the contents of the configuration file created by NetworkManager.

Resources:	
Files:	<code>/etc/sysconfig/network-scripts/ifcfg-eno1</code>
Machines:	<code>serverX</code>

Outcomes:

The **eno1** network interface on your **serverX** machine will be managed by NetworkManager with a connection named **eno1**. It will statically configure an IPv4 address of 192.168.0.1/24 without a gateway. The host with address 192.168.0.254 can be referenced as "otherhost".

Before you begin...

- Reset the **serverX** system.
- Log into and set up your **serverX** system. The following command makes the unconfigured **eno1** network interface available after the **serverX** system has been reset. It is used for all of the hands-on practice exercises and labs in the IPv6 chapter.

```
[student@serverX ~]$ lab ipv6 setup
```

- Become the **root** user.

```
[student@serverX ~]$ sudo -i
```

- 1. Before making any changes, display the list of existing network interfaces to determine the system's starting configuration. Also determine which interfaces are managed by NetworkManager.
 - 1.1. The **ip link** command will display all of the network interfaces recognized by the system.

```
[root@serverX ~]# ip link
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode
  DEFAULT
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state
  UP mode DEFAULT qlen 1000
    link/ether 52:54:00:00:07:0b brd ff:ff:ff:ff:ff:ff
4: eno1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state
  UP mode DEFAULT qlen 1000
    link/ether ca:8a:8f:84:e4:8f brd ff:ff:ff:ff:ff:ff
```

- 1.2. Use the **nmcli** to list the network interfaces that NetworkManager manages.

```
[root@serverX ~]# nmcli con show
```

NAME	UUID	TYPE	DEVICE
System eth0	5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03	802-3-ethernet	eth0

- 2. Create a NetworkManager connection, called **eno1**, for the **eno1** network interface. Redisplay the list of managed interfaces to confirm NetworkManager manages **eno1**.
- 2.1. Use **nmcli** to create the connection for **eno1**.

```
[root@serverX ~]# nmcli con add con-name eno1 type ethernet ifname eno1
Connection 'eno1' (6e2fa636-f7db-4567-bcba-6d12d0bbcc49) successfully
added.
```

- 2.2. Display the new list of interfaces managed by NetworkManager. **eno1** should be somewhere in the list.

```
[root@serverX ~]# nmcli con show
```

NAME	UUID	TYPE	DEVICE
eno1	6e2fa636-f7db-4567-bcba-6d12d0bbcc49	802-3-ethernet	eno1
System eth0	5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03	802-3-ethernet	eth0

- 3. Display the current IP address information for **eno1**.

```
[root@serverX ~]# ip addr show eno1
4: eno1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
qlen 1000
    link/ether ca:8a:8f:84:e4:8f brd ff:ff:ff:ff:ff:ff
    inet6 fe80::c88a:8fff:fe84:e48f/64 scope link
        valid_lft forever preferred_lft forever
```

It will have an IPv6 link-local address assigned to it (the address starting with **fe80::**), but it will not have an automatic IPv4 address.

- 4. Display the initial, default NetworkManager IPv4 configuration settings for the connection.

```
[root@serverX ~]# nmcli con show eno1 | grep ipv4
ipv4.method:                auto
ipv4.dns:
ipv4.dns-search:
ipv4.addresses:
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
```

- 5. Configure **eno1** to have a static IPv4 address of 192.168.0.1/24, without an additional gateway.

```
[root@serverX ~]# nmcli con mod eno1 ipv4.addresses '192.168.0.1/24'
[root@serverX ~]# nmcli con mod eno1 ipv4.method manual
```

- 6. Restart the **eno1** network interface and confirm its new IPv4 address configuration.
- 6.1. Bounce the **eno1** interface by taking it down, then bringing it back up.

```
[root@serverX ~]# nmcli con down eno1
[root@serverX ~]# nmcli con up eno1
Connection successfully activated (D-Bus active path: /org/freedesktop/
NetworkManager/ActiveConnection/2)
```

- 6.2. Use the **ip addr** command to confirm the interface's configuration.

```
[root@serverX ~]# ip addr show dev eno1
4: eno1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state
UP qlen 1000
    link/ether ca:8a:8f:84:e4:8f brd ff:ff:ff:ff:ff:ff
    inet 192.168.0.1/24 brd 192.168.0.255 scope global eno1
        valid_lft forever preferred_lft forever
    inet6 fe80::c88a:8fff:fe84:e48f/64 scope link
        valid_lft forever preferred_lft forever
```

Notice the new inet address entry.

- 7. Ping **eno1**'s own IPv4 address.

```
[root@serverX ~]# ping 192.168.0.1
PING 192.168.0.1 (192.168.0.1) 56(84) bytes of data.
64 bytes from 192.168.0.1: icmp_seq=1 ttl=64 time=0.091 ms
64 bytes from 192.168.0.1: icmp_seq=2 ttl=64 time=0.043 ms
^C
--- 192.168.0.1 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1006ms
rtt min/avg/max/mdev = 0.043/0.067/0.091/0.024 ms
```

- 8. Ping another IPv4 host, with address 192.168.0.254, to make sure it is reachable by **eno1**.

```
[root@serverX ~]# ping 192.168.0.254
PING 192.168.0.254 (192.168.0.254) 56(84) bytes of data.
64 bytes from 192.168.0.254: icmp_seq=1 ttl=64 time=0.165 ms
64 bytes from 192.168.0.254: icmp_seq=2 ttl=64 time=0.082 ms
^C
--- 192.168.0.254 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1010ms
rtt min/avg/max/mdev = 0.082/0.123/0.165/0.042 ms
```

- 9. Use the **ip** command to display the IPv4 routing table.


```
[root@serverX ~]# ip route
default via 172.25.7.254 dev eth0 proto static metric 1024
172.25.7.0/24 dev eth0 proto kernel scope link src 172.25.7.11
172.25.253.254 via 172.25.7.254 dev eth0 proto static metric 1
192.168.0.0/24 dev eno1 proto kernel scope link src 192.168.0.1
```

- 10. Identify the interface configuration file for the **eno1** network interface in **/etc/sysconfig/network-scripts**. View the file contents and note which variable assignments relate to the IPv4 configuration that was performed earlier.

```
[root@serverX ~]# ls /etc/sysconfig/network-scripts/ifcfg-*
/etc/sysconfig/network-scripts/ifcfg-eno1
/etc/sysconfig/network-scripts/ifcfg-eth0
/etc/sysconfig/network-scripts/ifcfg-lo
[root@serverX ~]# cat /etc/sysconfig/network-scripts/ifcfg-eno1
TYPE=Ethernet
BOOTPROTO=none
DEFROUTE=yes
IPV4_FAILURE_FATAL=no
IPV6INIT=yes
IPV6_AUTOCONF=yes
IPV6_DEFROUTE=yes
IPV6_FAILURE_FATAL=no
NAME=eno1
UUID=6e2fa636-f7db-4567-bcba-6d12d0bbcc49
DEVICE=eno1
ONBOOT=yes
IPADDR=192.168.0.1
PREFIX=24
IPV6_PEERDNS=yes
IPV6_PEERROUTES=yes
```

- 11. Configure the **hosts** file so that 192.168.0.254 can be referenced as "otherhost".

```
[root@serverX ~]# echo '192.168.0.254 otherhost' >> /etc/hosts
```

- 12. Ping "otherhost" by name to make sure the alias works.

```
[root@serverX ~]# ping otherhost
PING otherhost (192.168.0.254) 56(84) bytes of data.
64 bytes from otherhost (192.168.0.254): icmp_seq=1 ttl=64 time=0.099 ms
64 bytes from otherhost (192.168.0.254): icmp_seq=2 ttl=64 time=0.070 ms
^C
--- otherhost ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1009ms
rtt min/avg/max/mdev = 0.070/0.084/0.099/0.017 ms
```

IPv6 Networking Concepts

Objectives

After completing this section, students should be able to explain the basic concepts of IPv6 addresses and networking.

IPv6 overview

IPv6 is intended as the replacement for the IPv4 network protocol. The major problem it solves is the exhaustion of IPv4 addresses by using a much larger network address space. It also provides a number of enhancements and new features for network configuration management and support for future protocol changes.

The key reason IPv6 is not yet in wide deployment is that the core protocol does not have a simple way for systems that only have IPv6 addresses to communicate with systems that only have IPv4 addresses.

The best transition plan at present is to provide all hosts with both IPv4 and IPv6 addresses, so that Internet resources only using one of the protocols can be reached from the host. This is called a *dual-stack* configuration, and is the approach on which this course will focus.



Note

There are a number of promising transition methods in development to allow IPv6-only hosts to use the IPv4 Internet or support other forms of IPv4/IPv6 translation, such as NAT64 (RFC 6145) and 464XLAT (RFC 6877), but they are beyond the scope of this course.

The basic position of the Internet Engineering Task Force (IETF) is that network operators using IPv4 should “obtain an IPv6 prefix, turn on IPv6 routing within their networks and between themselves and any peer, upstream, or downstream neighbors, enable it on their computers, and use it in normal processing. This should be done while leaving IPv4 stable, until a point is reached that any communication that can be carried out could use either protocol equally well. At that point, the economic justification for running both becomes debatable, and network operators can justifiably turn IPv4 off.” (RFC 6144, Introduction)

Interpreting IPv6 addresses

IPv6 addresses

An IPv6 address is a 128-bit number, normally expressed as eight colon-separated groups of four hexadecimal nibbles (half-bytes). Each nibble represents four bits of the IPv6 address, so each group represents 16 bits of the IPv6 address.

```
2001:0db8:0000:0010:0000:0000:0000:0001
```

To make it easier to write IPv6 addresses, leading zeros in a colon-separated group do not need to be written. However, at least one nibble must be written in each field. Zeros which follow a nonzero nibble in the group do need to be written.

```
2001:db8:0:10:0:0:0:1
```

Since addresses with long strings of zeros are common, one or more groups of consecutive zeros may be combined with *exactly one* `::` block.

```
2001:db8:0:10::1
```

Notice that under these rules, `2001:db8::0010:0:0:0:1` would be another less convenient way to write the example address. But it is a valid representation of the same address, and this can confuse administrators new to IPv6. Some tips for writing consistently readable addresses:

1. Leading zeros in a group must always be suppressed.
2. Use `::` to shorten as much as possible. If two runs of zeros are equal in length, shorten the leftmost run of zeros by preference.
3. Although it is allowed, do not use `::` to shorten one group of zeros. Use `:0:` instead, and save `::` for runs of zeros longer than a single group.
4. Always use lowercase letters for hexadecimal numbers **a** through **f**.



Important

When including a TCP or UDP network port after an IPv6 address, always enclose the IPv6 address in square brackets so that the port does not look like it is part of the address.

```
[2001:db8:0:10::1]:80
```

IPv6 subnets

A normal unicast address is divided into two parts: the *network prefix* and *interface ID*. The network prefix identifies the subnet. No two network interfaces on the same subnet can have the same interface ID; the interface ID identifies a particular interface on the subnet.

Unlike IPv4, IPv6 has a standard subnet mask, which is used for almost all normal addresses, **/64**. In this case, half of the address is the network prefix and half of it is the interface ID. This means that a single subnet can hold as many hosts as necessary.

Typically, the network provider will allocate a shorter prefix to an organization, such as a **/48**. This leaves the rest of the network part for assigning subnets from that allocated prefix. For a **/48** allocation, that leaves 16 bits for subnets (up to 65536 subnets).

IPv6 address is **2001:db8:0:1::1/64**
 Allocation from provider is **2001:db8::/48**




Figure 2.1: IPv6 address parts and subnetting

IPv6 address allocation

Common IPv6 Addresses and Networks

IPv6 address or network	Purpose	Description
::1/128	localhost	The IPv6 equivalent to 127.0.0.1/8 , set on the loopback interface.
::	The unspecified address	The IPv6 equivalent to 0.0.0.0 . For a network service, this could indicate that it is listening on all configured IP addresses.
::/0	The default route (the IPv6 Internet)	The IPv6 equivalent to 0.0.0.0/0 . The default route in the routing table matches this network; the router for this network is where all traffic is sent for which there is not a better route.
2000::/3	Global unicast addresses	"Normal" IPv6 addresses are currently being allocated from this space by IANA. This is equivalent to all the networks ranging from 2000::/16 through 3fff::/16 .
fd00::/8	Unique local addresses (RFC 4193)	IPv6 has no direct equivalent of RFC 1918 private address space, although this is close. A site can use these to self-allocate a private routable IP address space inside the organization, but these networks cannot be used on the global Internet. The site must <i>randomly</i> select a /48 from this space, but it can subnet the allocation into /64 networks normally.
fe80::/64	Link-local addresses	Every IPv6 interface automatically configures a <i>link-local</i> address that only works on the local link on this network. This will be discussed in more detail later.

IPv6 address or network	Purpose	Description
ff00::/8	Multicast	The IPv6 equivalent to 224.0.0.0/4 . Multicast is used to transmit to multiple hosts at the same time, and is particularly important in IPv6 because it has no broadcast addresses.

Link-local addresses

A *link-local address* in IPv6 is an unroutable address which is used only to talk to hosts on a specific network link. Every network interface on the system is automatically configured with a link-local address on the **fe80::** network. To ensure that it is unique, the interface ID of the link-local address is constructed from the network interface's Ethernet hardware address. The usual procedure to convert the 48-bit MAC address to a 64-bit interface ID is to set bit 7 of the MAC address and insert **ff:fe** between its two middle bytes.

- Network prefix: **fe80::/64**
- MAC address: **00:11:22:aa:bb:cc**
- Link-local address: **fe80::211:22ff:feaa:bbcc/64**

The link-local addresses of other machines can be used like normal addresses by other hosts on the same link. Since *every* link has a **fe80::/64** network on it, the routing table cannot be used to select the outbound interface correctly. The link to use when talking to a link-local address must be specified with a *scope identifier* at the end of the address. The scope identifier consists of % followed by the name of the network interface.

For example, to use **ping6** to ping the link-local address **fe80::211:22ff:feaa:bbcc** using the link connected to the **eth0** network interface, the correct command would be:

```
[student@demo ~]$ ping6 fe80::211:22ff:feaa:bbcc%eth0
```



Note

Scope identifiers are only needed when contacting addresses that have "link" scope. Normal global addresses are used just like they are in IPv4, and select their outbound interfaces from the routing table.

Multicast

Multicast plays a larger role in IPv6 than in IPv4 because there is no broadcast address in IPv6. One key multicast address in IPv6 is **ff02::1**, the all-nodes link-local address. Pinging this address will send traffic to all nodes on the link. Link-scope multicast addresses (starting **ff02::/8**) need to be specified with a scope identifier, just like a link-local address.

```
[student@demo ~]$ ping6 ff02::1%eth0
PING ff02::1%eth0(ff02::1) 56 data bytes
64 bytes from fe80::211:22ff:feaa:bbcc: icmp_seq=1 ttl=64 time=0.072 ms
64 bytes from fe80::200:aaff:fe33:2211: icmp_seq=1 ttl=64 time=102 ms (DUP!)
64 bytes from fe80::bcd:eff:fea1:b2c3: icmp_seq=1 ttl=64 time=103 ms (DUP!)
64 bytes from fe80::211:22ff:feaa:bbcc: icmp_seq=2 ttl=64 time=0.079 ms
...
```

IPv6 address configuration

IPv4 has two ways in which addresses get configured on network interfaces. Network addresses may be configured on interfaces manually by the administrator, or dynamically from the network using DHCP. IPv6 also supports manual configuration, and two methods of dynamic configuration, one of which is DHCPv6.

Static addressing

Interface IDs for static IPv6 addresses can be selected at will, just like IPv4. In IPv4, there were two addresses on a network that could not be used, the lowest address in the subnet and the highest address in the subnet. In IPv6, the following interface IDs are reserved and cannot be used for a normal network address on a host.

- The all-zeros identifier **0000:0000:0000:0000** ("subnet router anycast") used by all routers on the link. (For the **2001:db8::/64** network, this would be the address **2001:db8::**.)
- The identifiers **fdff:ffff:ffff:ff80** through **fdff:ffff:ffff:ffff**.

DHCPv6 configuration

DHCPv6 works a little differently than DHCP for IPv4, because there is no broadcast address. Essentially, a host sends a DHCPv6 request from its link-local address to port 547/UDP on **ff02::1:2**, the all-dhcp-servers link-local multicast group. The DHCPv6 server then usually sends a reply with appropriate information to port 546/UDP on the client's link-local address.

The *dhcp* package in RHEL 7 provides support for a DHCPv6 server.

SLAAC configuration

In addition to DHCPv6, IPv6 also supports a second dynamic configuration method, called *Stateless Address Autoconfiguration* (SLAAC). Using SLAAC, the host brings up its interface with a link-local **fe80::/64** address normally. It then sends a "router solicitation" to **ff02::2**, the all-routers link-local multicast group. An IPv6 router on the local link responds to the host's link-local address with a network prefix and possibly other information. The host then uses that network prefix with an interface ID that it normally constructs in the same way that link-local addresses are constructed. The router periodically sends multicast updates ("router advertisements") to confirm or update the information it provided.

The *radvd* package in RHEL 7 allows a RHEL-based IPv6 router to provide SLAAC through router advertisements.



Important

A typical RHEL 7 machine configured to get IPv4 addresses through DHCP is usually also configured to use SLAAC to get IPv6 addresses. This can result in machines unexpectedly obtaining IPv6 addresses when an IPv6 router is added to the network.

Some IPv6 deployments combine SLAAC and DHCPv6, using SLAAC to only provide network address information and DHCPv6 to provide other information, such as which DNS servers and search domains to configure.



References

ping(8), **radvd(8)**, and **dhcpcd(8)** man pages

Selected IETF RFC references:

RFC 2460: Internet Protocol, Version 6 (IPv6) Specification
<http://tools.ietf.org/html/rfc2460>

RFC 4291: IP Version 6 Addressing Architecture
<http://tools.ietf.org/html/rfc4291>

RFC 5952: A Recommendation For IPv6 Address Text Representation
<http://tools.ietf.org/html/rfc5952>

RFC 4862: IPv6 Stateless Address Autoconfiguration
<http://tools.ietf.org/html/rfc4862>

RFC 3315: Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
<http://tools.ietf.org/html/rfc3315>

RFC 3736: Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6
<http://tools.ietf.org/html/rfc3736>

RFC 4193: Unique Local IPv6 Unicast Addresses
<http://tools.ietf.org/html/rfc4193>

Practice: Interpreting IPv6 Addresses

Quiz

Match the following compressed IPv6 addresses to their counterparts in the table.

2000::1	2001:3:700::2	2001:3::7:0:2	
2001:db8:0:7::2	2001:db8::7::2	2::1	::
ff02::1:0:0			

IPv6 address	Compressed IPv6 address
2000:0000:0000:0000:0000:0000:0000:0001	
0002:0000:0000:0000:0000:0000:0000:0001	
2001:0db8:0000:0007:0000:0000:0000:0002	
2001:0003:0000:0000:0000:0007:0000:0002	
2001:0003:0700:0000:0000:0000:0000:0002	
ff02:0000:0000:0000:0000:0001:0000:0000	
0000:0000:0000:0000:0000:0000:0000:0000	
Not a valid IPv6 address	

Solution

Match the following compressed IPv6 addresses to their counterparts in the table.

IPv6 address	Compressed IPv6 address
2000:0000:0000:0000:0000:0000:0000:0001	2000::1
0002:0000:0000:0000:0000:0000:0000:0001	2::1
2001:0db8:0000:0007:0000:0000:0000:0002	2001:db8:0:7::2
2001:0003:0000:0000:0000:0007:0000:0002	2001:3::7:0:2
2001:0003:0700:0000:0000:0000:0000:0002	2001:3:700::2
ff02:0000:0000:0000:0000:0001:0000:0000	ff02::1:0:0
0000:0000:0000:0000:0000:0000:0000:0000	::
Not a valid IPv6 address	2001:db8::7::2

IPv6 Networking Configuration

Objectives

After completing this section, students should be able to configure IPv6 networking using **nmcli** and configuration files in the **/etc/sysconfig/network-scripts** directory.

NetworkManager and IPv6

To work with IPv6 addresses using NetworkManager, all the commands that are used with IPv4 networking work with IPv6 networking. There are some different settings that are relevant for connections, but most commands will be similar for IPv6 configuration.

Adding an IPv6 network connection

The **nmcli con add** command is used to add new network connections.

The following command, shown in a previous section of this chapter, will add a new connection for the interface **eno2**, which will autoconnect at startup, getting IPv4 networking information using DHCPv4. It will also get IPv6 networking settings by listening for router advertisements on the local link.

```
[root@demo ~]# nmcli con add con-name eno2 type ethernet ifname eno2
```

The next example configures the **eno2** interface statically instead, using the IPv6 address and network prefix **2001:db8:0:1::c000:207/64** and default IPv6 gateway **2001:db8:0:1::1**, and the IPv4 address and network prefix **192.0.2.7/24** and default IPv4 gateway **192.0.2.1**, but still autoconnects at startup and saves its configuration into **/etc/sysconfig/network-scripts/ifcfg-eno2**. The example is line-wrapped with a shell **** escape.

```
[root@demo ~]# nmcli con add con-name eno2 type ethernet ifname eno2 \  
> ip6 2001:db8:0:1::c000:207/64 gw6 2001:db8:0:1::1 ip4 192.0.2.7/24 gw4 192.0.2.1
```

Modifying network connection settings for IPv6

The **nmcli con show *name*** command, where *name* is the name of the connection, can be used to view IPv6-related settings:

```
[root@demo ~]# nmcli con show static-eth0 | grep ipv6  
ipv6.method:                manual  
ipv6.dns:                    2001:4860:4860::8888  
ipv6.dns-search:             example.com  
ipv6.addresses:              { ip = 2001:db8:0:1::7/64, gw = 2001:db8:0:1::1 }  
ipv6.routes:                   
ipv6.ignore-auto-routes:     no  
ipv6.ignore-auto-dns:        no  
ipv6.never-default:          no  
ipv6.may-fail:                yes  
ipv6.ip6-privacy:             -1 (unknown)  
ipv6.dhcp-hostname:          --  
[root@demo ~]#
```

Likewise, **nmcli con mod name** can be used to adjust how connections set IPv6 addresses.

To set the IPv6 address to 2001:db8:0:1::a00:1/64 and default gateway to 2001:db8:0:1::1 for the connection **static-eth0**:

```
[root@demo ~]# nmcli con mod static-eth0 ipv6.address "2001:db8:0:1::a00:1/64
2001:db8:0:1::1"
```



Important

If a connection that got its IPv6 information by SLAAC or a DHCPv6 server is being changed to get it from static configuration files only, the setting **ipv6.method** should also be changed from **auto** or **dhcp** to **manual**. Otherwise, the connection may hang or not complete successfully when it is activated, or it may get an IPv6 address from SLAAC or DHCPv6 in addition to the static address.

A number of settings may have multiple values. A specific value can be added to the list or deleted from the list for a setting by adding a **+** or **-** symbol to the start of the setting name.

To add the DNS server 2001:4860:4860::8888 to the list of nameservers to use with the connection **static-eth0**:

```
[root@demo ~]# nmcli con mod static-eth0 +ipv6.dns 2001:4860:4860::8888
```



Note

Static IPv4 and IPv6 DNS settings all end up as **nameserver** directives in **/etc/resolv.conf**. It may be a good idea to ensure that there is, at minimum, an IPv4-reachable nameserver (assuming a dual-stack system) and preferably at least one nameserver using each protocol in case of connectivity issues with either IPv4 or IPv6 networking.

Remember that the file **/etc/sysconfig/network-scripts/ifcfg-name** can be directly edited, and that **nmcli con reload** must be run after saving so that NetworkManager reads the configuration changes.

The following table maps some of the key NetworkManager setting names relevant to IPv6 connections to **ifcfg-*** directives.

Comparison of nm-settings and ifcfg-* Directives

nmcli con mod	ifcfg-* file	Effect
ipv6.method manual	IPV6_AUTOCONF=no	IPv6 addresses configured statically.
ipv6.method auto	IPV6_AUTOCONF=yes	Will configure network settings using SLAAC from router advertisements.
ipv6.method dhcp	IPV6_AUTOCONF=no DHCPV6C=yes	Will configure network settings by using DHCPv6, but not SLAAC.

nmcli con mod	ifcfg-* file	Effect
ipv6.addresses "2001:db8::a/64 2001:db8::1"	IPV6ADDR=2001:db8::a/64 IPV6_DEFAULTGW=2001:db8::1	Sets static IPv4 address, network prefix, and default gateway. If more than one address is set for the connection, IPV6_SECONDARIES takes a double-quoted list of space-delimited address/prefix definitions.
ipv6.dns . . .	DNS0= . . .	Modify /etc/resolv.conf to use this nameserver . Exactly the same as IPv4.
ipv6.dns-search example.com	DOMAIN=example.com	Modify /etc/resolv.conf to use this domain in the search directive. Exactly the same as IPv4.
ipv6.ignore-auto-dns true	IPV6_PEERDNS=no	Ignore DNS server information from the DHCP server.
connection.autoconnect yes	ONBOOT=yes	Automatically activate this connection at boot.
connection.id eth0	NAME=eth0	The name of this connection.
connection.interface-name eth0	DEVICE=eth0	The connection is bound to the network interface with this name.
802-3-ethernet.mac-address . . .	HWADDR= . . .	The connection is bound to the network interface with this MAC address.

Viewing IPv6 networking information

Both **nmcli dev status** to show the NetworkManager status of all network devices and **nmcli con show** to show the list of available connections work exactly as they do for IPv4.

The **ip addr show** command still displays the current configuration of network interfaces on the system. The example that follows calls out some items relevant to IPv6.

```
[student@demo ~]$ ip addr show eth0
2: eth0: <BROADCAST,MULTICAST, ①UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
    ②link/ether 52:54:00:00:00:0b brd ff:ff:ff:ff:ff:ff
    ③inet 192.0.2.2/24 brd 192.0.2.255 scope global eth0
        valid_lft forever preferred_lft forever
    ④inet6 2001:db8:0:1:5054:ff:fe00:b/64 scope global
        valid_lft forever preferred_lft forever
    ⑤inet6 fe80::5054:ff:fe00:b/64 scope link
        valid_lft forever preferred_lft forever
```

- ❶ An active interface is **UP**.
- ❷ The **link/ether** line specifies the hardware (MAC) address of the device.
- ❸ The **inet** line shows an IPv4 address, its network prefix length, and scope.
- ❹ The **inet6** line shows an IPv6 address, its network prefix length, and scope. This address is of *global* scope and is normally used.
- ❺ This **inet6** line is for an address of *link* scope and can only be used for communication on the local Ethernet link.

The **ip -6 route show** command displays the IPv6 routing table for the system:

```
[root@demo ~]# ip -6 route show
unreachable ::/96 dev lo metric 1024 error -101
unreachable ::ffff:0.0.0.0/96 dev lo metric 1024 error -101
2001:db8:0:1::/64 dev eth0 proto kernel metric 256
unreachable 2002:a00::/24 dev lo metric 1024 error -101
unreachable 2002:7f00::/24 dev lo metric 1024 error -101
unreachable 2002:a9fe::/32 dev lo metric 1024 error -101
unreachable 2002:ac10::/28 dev lo metric 1024 error -101
unreachable 2002:c0a8::/32 dev lo metric 1024 error -101
unreachable 2002:e000::/19 dev lo metric 1024 error -101
unreachable 3ffe:ffff::/32 dev lo metric 1024 error -101
fe80::/64 dev eth0 proto kernel metric 256
default via 2001:db8:0:1::ffff dev eth0 proto static metric 1024
```

In the previous example, ignore the **unreachable** routes, which point at networks which are never to be used. That leaves three routes:

1. To the 2001:db8:0:1::/64 network using the eth0 interface (which presumably has an address on that network).
2. To the fe80::/64 network using the eth0 interface, for the link-local address. On a system with multiple interfaces, there will be a route to fe80::/64 out each interface for each link-local address.
3. A default route to all networks on the IPv6 Internet (the ::/0 network) that don't have a more specific route on the system, through the router at 2001:db8:0:1::ffff, reachable with the eth0 device.

IPv6 troubleshooting tools

Connectivity

The **ping6** command is the IPv6 version of **ping** in Red Hat Enterprise Linux. It communicates over IPv6 and can take IPv6 addresses, but otherwise works like **ping**.

```
[root@demo ~]# ping6 2001:db8:0:1::1
PING 2001:db8:0:1::1(2001:db8:0:1::1) 56 data bytes
64 bytes from 2001:db8:0:1::1: icmp_seq=1 ttl=64 time=18.4 ms
64 bytes from 2001:db8:0:1::1: icmp_seq=2 ttl=64 time=0.178 ms
64 bytes from 2001:db8:0:1::1: icmp_seq=3 ttl=64 time=0.180 ms
^C
--- 2001:db8:0:1::1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2001ms
rtt min/avg/max/mdev = 0.178/6.272/18.458/8.616 ms
[root@demo ~]#
```

Link-local addresses and the link-local all-nodes multicast group (ff02::1) can be pinged, but the network interface to use must be specified explicitly with a scope zone identifier (such as ff02::1%eth0). If this is left out, the error **connect: Invalid argument** will be displayed.

Pinging ff02::1 can be useful for finding other IPv6 nodes on the local network.

```
[root@rhel7 ~]# ping6 ff02::1%eth1
PING ff02::1%eth1(ff02::1) 56 data bytes
64 bytes from fe80::78cf:7fff:fed2:f97b: icmp_seq=1 ttl=64 time=22.7 ms
64 bytes from fe80::f482:dbff:fe25:6a9f: icmp_seq=1 ttl=64 time=30.1 ms (DUP!)
64 bytes from fe80::78cf:7fff:fed2:f97b: icmp_seq=2 ttl=64 time=0.183 ms
64 bytes from fe80::f482:dbff:fe25:6a9f: icmp_seq=2 ttl=64 time=0.231 ms (DUP!)
^C
--- ff02::1%eth1 ping statistics ---
2 packets transmitted, 2 received, +2 duplicates, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 0.183/13.320/30.158/13.374 ms
[root@rhel7 ~]# ping6 -c 1 fe80::f482:dbff:fe25:6a9f%eth1
PING fe80::f482:dbff:fe25:6a9f%eth1(fe80::f482:dbff:fe25:6a9f) 56 data bytes
64 bytes from fe80::f482:dbff:fe25:6a9f: icmp_seq=1 ttl=64 time=22.9 ms

--- fe80::f482:dbff:fe25:6a9f%eth1 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 22.903/22.903/22.903/0.000 ms
```

Remember that IPv6 link-local addresses can be used by other hosts on the same link, just like normal addresses.

```
[student@demo ~]$ ssh fe80::f482:dbff:fe25:6a9f%eth1
student@fe80::f482:dbff:fe25:6a9f%eth1's password:
Last login: Thu Jun  5 15:20:10 2014 from demo.example.com
[student@server ~]$
```

Routing

The **tracepath6** and **traceroute -6** commands are the equivalent to **tracepath** and **traceroute** for IPv6.

```
[root@demo ~]# tracepath6 2001:db8:0:2::451
1?: [LOCALHOST] 0.091ms pmtu 1500
1: 2001:db8:0:1::ba 0.214ms
2: 2001:db8:0:1::1 0.512ms
3: 2001:db8:0:2::451 0.559ms reached
Resume: pmtu 1500 hops 3 back 3
```

Ports and services

Either the **ss** command or the **netstat** command can display information about network sockets, and they take almost identical options.

```
[root@demo ~]# ss -A inet -n
Netid State Recv-Q Send-Q Local Address:Port Peer Address:Port
tcp ESTAB 0 0 192.168.122.98:22 192.168.122.1:35279
tcp ESTAB 0 0 2001:db8:0:1::ba:22 2001:db8:0:1::1:40810
[root@demo ~]# netstat -46n
Active Internet connections (w/o servers)
Proto Recv-Q Send-Q Local Address Foreign Address State
tcp 0 0 192.168.122.98:22 192.168.122.1:35279 ESTABLISHED
tcp6 0 0 2001:db8:0:1::ba:22 2001:db8:0:1::1:40810 ESTABLISHED
```

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.
-A inet	Display active connections (but not listening sockets) for the inet address family. That is, ignore local UNIX domain sockets. For ss , both IPv4 and IPv6 connections will be displayed. For netstat , only IPv4 connections will be displayed. (netstat -A inet6 will display IPv6 connections, and netstat -46 will display IPv4 and IPv6 at the same time.)

R

References

NetworkManager(8), **nmcli**(1), **nmcli-examples**(5), **nm-settings**(5), **ip**(8), **ip-address**(8), **ip-route**(8), **ping6**(8), **tracepath6**(8), **traceroute**(8), **ss**(8), and **netstat**(8) man pages

Practice: Configuring IPv6 Networking

Guided exercise

In this lab, you will configure a network interface with a static IPv6 address. Once the interface is configured, you will confirm that it works and identify other IPv6 nodes on the local network. You will also explore the contents of the configuration file created by NetworkManager.

Resources:	
Files:	<code>/etc/sysconfig/network-scripts/ifcfg-eno1</code>
Machines:	<code>serverX</code>

Outcomes:

The **eno1** network interface on your **serverX** machine will be managed by NetworkManager with a connection named **eno1**. It will statically configure an IPv6 address of `fddb:fe2a:ab1e::c0a8:1/64` and use `fddb:fe2a:ab1e::c0a8:fe/64` as the IPv6 gateway.

Before you begin...

- Reset the **serverX** system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab ipv6 setup
```

- Become the **root** user.

```
[student@serverX ~]$ sudo -i
```

- 1. Before making any changes, display the list of existing network interfaces in order to determine the system's current configuration. Also determine which interfaces are managed by NetworkManager.
 - 1.1. The **ip link** command will display all of the network interfaces recognized by the system.

```
[root@serverX ~]# ip link
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode
  DEFAULT
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state
  UP mode DEFAULT qlen 1000
    link/ether 52:54:00:00:07:0b brd ff:ff:ff:ff:ff:ff
4: eno1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state
  UP mode DEFAULT qlen 1000
    link/ether ce:c4:7c:28:4c:7a brd ff:ff:ff:ff:ff:ff
```

- 1.2. Use the **nmcli** to list the network interfaces that NetworkManager manages.


```
[root@serverX ~]# nmcli con show
NAME                UUID                                TYPE                DEVICE
System eth0        5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03  802-3-ethernet     eth0
```

- 2. Create a NetworkManager connection, called **eno1**, for the **eno1** network interface. Redisplay the list of managed interfaces to confirm NetworkManager manages **eno1**.
 - 2.1. Use **nmcli** to create the connection for **eno1**.

```
[root@serverX ~]# nmcli con add con-name eno1 type ethernet ifname eno1
Connection 'eno1' (0d687259-c64b-4e5b-bece-cabbe952e46f) successfully
added.
```

- 2.2. Display the new list of interfaces managed by NetworkManager. **eno1** should be somewhere in the list.

```
[root@serverX ~]# nmcli con show
NAME                UUID                                TYPE                DEVICE
eno1                0d687259-c64b-4e5b-bece-cabbe952e46f  802-3-ethernet     eno1
System eth0        5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03  802-3-ethernet     eth0
```

- 3. Display the current IP address information for **eno1**.

```
[root@serverX ~]# ip addr show eno1
4: eno1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
qlen 1000
    link/ether 06:8f:6e:13:6e:8e brd ff:ff:ff:ff:ff:ff
    inet6 fe80::48f:6eff:fe13:6e8e/64 scope link
        valid_lft forever preferred_lft forever
```

It will have an IPv6 link-local address assigned to it (the address on the **fe80::/64** network).

- 4. Display the initial, default NetworkManager IPv6 configuration settings for the connection.

```
[root@serverX ~]# nmcli con show eno1 | grep ipv6
ipv6.method:                auto
ipv6.dns:
ipv6.dns-search:
ipv6.addresses:
ipv6.routes:
ipv6.ignore-auto-routes:    no
ipv6.ignore-auto-dns:      no
ipv6.never-default:        no
ipv6.may-fail:              yes
ipv6.ip6-privacy:           -1 (unknown)
ipv6.dhcp-hostname:        --
```

- 5. Configure **eno1** to have a static IPv6 address of **fddb:fe2a:ab1e::c0a8:1** with a standard **/64** subnet prefix. Use **fddb:fe2a:ab1e::c0a8:fe** as the IPv6 gateway.

```
[root@serverX ~]# nmcli con mod eno1 ipv6.addresses 'fddb:fe2a:ab1e::c0a8:1/64
fddb:fe2a:ab1e::c0a8:fe'
[root@serverX ~]# nmcli con mod eno1 ipv6.method manual
```

- 6. Restart the **eno1** network interface and confirm its new IPv6 address configuration.
 - 6.1. Bounce the **eno1** interface by taking it down, then bringing it back up.

```
[root@serverX ~]# nmcli con down eno1
[root@serverX ~]# nmcli con up eno1
Connection successfully activated (D-Bus active path: /org/freedesktop/
NetworkManager/ActiveConnection/2)
```

- 6.2. Use the **ip addr** command to confirm the interface's configuration.

```
[root@serverX ~]# ip addr show dev eno1
4: eno1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state
UP qlen 1000
    link/ether 06:8f:6e:13:6e:8e brd ff:ff:ff:ff:ff:ff
    inet6 fddb:fe2a:ab1e::c0a8:1/64 scope global
        valid_lft forever preferred_lft forever
    inet6 fe80::48f:6eff:fe13:6e8e/64 scope link
        valid_lft forever preferred_lft forever
```

Notice the global address, fddb:fe2a:ab1e::c0a8:1/64, is available for use.

- 7. Ping **eno1**'s own IPv6 address.

```
[root@serverX ~]# ping6 fddb:fe2a:ab1e::c0a8:1
PING fddb:fe2a:ab1e::c0a8:1(fddb:fe2a:ab1e::c0a8:1) 56 data bytes
64 bytes from fddb:fe2a:ab1e::c0a8:1: icmp_seq=1 ttl=64 time=0.141 ms
64 bytes from fddb:fe2a:ab1e::c0a8:1: icmp_seq=2 ttl=64 time=0.081 ms
^C
--- fddb:fe2a:ab1e::c0a8:1 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 999ms
rtt min/avg/max/mdev = 0.081/0.111/0.141/0.030 ms
```

- 8. Ping the IPv6 gateway to make sure it is reachable by **eno1**.

```
[root@serverX ~]# ping6 fddb:fe2a:ab1e::c0a8:fe
PING fddb:fe2a:ab1e::c0a8:fe(fddb:fe2a:ab1e::c0a8:fe) 56 data bytes
64 bytes from fddb:fe2a:ab1e::c0a8:fe: icmp_seq=1 ttl=64 time=0.254 ms
64 bytes from fddb:fe2a:ab1e::c0a8:fe: icmp_seq=2 ttl=64 time=0.123 ms
64 bytes from fddb:fe2a:ab1e::c0a8:fe: icmp_seq=3 ttl=64 time=0.119 ms
64 bytes from fddb:fe2a:ab1e::c0a8:fe: icmp_seq=4 ttl=64 time=0.123 ms
64 bytes from fddb:fe2a:ab1e::c0a8:fe: icmp_seq=5 ttl=64 time=0.090 ms
^C
--- fddb:fe2a:ab1e::c0a8:fe ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4000ms
rtt min/avg/max/mdev = 0.090/0.141/0.254/0.059 ms
```

- 9. Use the **ip** command to display the IPv6 routing table. Note the default gateway listed for IPv6.

```
[root@serverX ~]# ip -6 route
unreachable ::/96 dev lo metric 1024 error -101
unreachable ::ffff:0.0.0.0/96 dev lo metric 1024 error -101
unreachable 2002:a00::/24 dev lo metric 1024 error -101
unreachable 2002:7f00::/24 dev lo metric 1024 error -101
unreachable 2002:a9fe::/32 dev lo metric 1024 error -101
unreachable 2002:ac10::/28 dev lo metric 1024 error -101
unreachable 2002:c0a8::/32 dev lo metric 1024 error -101
unreachable 2002:e000::/19 dev lo metric 1024 error -101
unreachable 3ffe:ffff::/32 dev lo metric 1024 error -101
fddb:fe2a:ab1e::/64 dev eno1 proto kernel metric 256
fe80::/64 dev eth0 proto kernel metric 256
fe80::/64 dev eno1 proto kernel metric 256
default via fddb:fe2a:ab1e::c0a8:fe dev eno1 proto static metric 1024
```

- 10. Discover other local IPv6 nodes on the network. Ping the link-local all-nodes multicast group (**ff02::1**) through the **eno1** interface to see what other hosts respond.

```
[root@serverX ~]# ping6 ff02::1%eno1
PING ff02::1%eno1(ff02::1) 56 data bytes
64 bytes from fe80::fc46:acff:feffe:10b7: icmp_seq=1 ttl=64 time=0.298 ms
64 bytes from fe80::707e:68ff:fe3e:fd23: icmp_seq=1 ttl=64 time=0.306 ms (DUP!)
64 bytes from fe80::707e:68ff:fe3e:fd23: icmp_seq=2 ttl=64 time=0.125 ms
64 bytes from fe80::fc46:acff:feffe:10b7: icmp_seq=2 ttl=64 time=0.161 ms (DUP!)
64 bytes from fe80::707e:68ff:fe3e:fd23: icmp_seq=3 ttl=64 time=0.107 ms
64 bytes from fe80::fc46:acff:feffe:10b7: icmp_seq=3 ttl=64 time=0.136 ms (DUP!)
64 bytes from fe80::707e:68ff:fe3e:fd23: icmp_seq=4 ttl=64 time=0.111 ms
64 bytes from fe80::fc46:acff:feffe:10b7: icmp_seq=4 ttl=64 time=0.143 ms (DUP!)
64 bytes from fe80::707e:68ff:fe3e:fd23: icmp_seq=5 ttl=64 time=0.131 ms
64 bytes from fe80::fc46:acff:feffe:10b7: icmp_seq=5 ttl=64 time=0.167 ms (DUP!)
64 bytes from fe80::707e:68ff:fe3e:fd23: icmp_seq=6 ttl=64 time=0.109 ms
64 bytes from fe80::fc46:acff:feffe:10b7: icmp_seq=6 ttl=64 time=0.141 ms (DUP!)
64 bytes from fe80::707e:68ff:fe3e:fd23: icmp_seq=7 ttl=64 time=0.116 ms
64 bytes from fe80::fc46:acff:feffe:10b7: icmp_seq=7 ttl=64 time=0.150 ms (DUP!)
^C
--- ff02::1%eno1 ping statistics ---
7 packets transmitted, 7 received, +7 duplicates, 0% packet loss, time 5999ms
rtt min/avg/max/mdev = 0.107/0.157/0.306/0.062 ms
```

- 11. Identify the interface configuration file for the **eno1** network interface in **/etc/sysconfig/network-scripts**. View the file contents and note which variable assignments relate to the IPv6 configuration that was performed earlier.

```
[root@serverX ~]# ls /etc/sysconfig/network-scripts/ifcfg-*
/etc/sysconfig/network-scripts/ifcfg-eno1
/etc/sysconfig/network-scripts/ifcfg-eth0
/etc/sysconfig/network-scripts/ifcfg-lo
[root@serverX ~]# cat /etc/sysconfig/network-scripts/ifcfg-eno1
TYPE=Ethernet
BOOTPROTO=dhcp
DEFROUTE=yes
IPV4_FAILURE_FATAL=no
IPV6INIT=yes
IPV6_AUTOCONF=no
IPV6_DEFROUTE=yes
IPV6_FAILURE_FATAL=no
NAME=eno1
UUID=4214d89b-f409-4853-8e31-4e673845e1a1
DEVICE=eno1
ONBOOT=yes
IPV6ADDR=fddb:fe2a:ab1e::c0a8:1/64
IPV6_DEFAULTGW=fddb:fe2a:ab1e::c0a8:fe
PEERDNS=yes
PEERROUTES=yes
```

Lab: Managing IPv6 Networking

Performance checklist

In this lab, you will assign static IPv4 and IPv6 addresses to a network interface.

Resources:

Machines:	serverX
------------------	----------------

Outcomes:

serverX will have a NetworkManager connection, called **eno1**, that manages the **eno1** network interface. It will be configured statically with an IPv4 address of 192.168.0.100/24 and an IPv6 address of fddb:fe2a:ab1e::c0a8:64/64.

Before you begin...

- Reset the serverX system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab ipv6 setup
```

- Become the **root** user.

```
[student@serverX ~]$ sudo -i
```

1. Create a NetworkManager connection, called **eno1**, that corresponds to the **eno1** network interface.
2. Configure **eno1** with a static IPv4 address of 192.168.0.100/24.
3. Configure **eno1** with a static IPv6 address of fddb:fe2a:ab1e::c0a8:64/64.
4. Restart the **eno1** network interface and confirm its new IPv4 and IPv6 address configuration.
5. Ping the local IPv4 gateway, 192.168.0.254, and the local IPv6 gateway, fddb:fe2a:ab1e::c0a8:fe, to confirm they can both be reached through the **eno1** interface.
6. Run the grading script to verify your work.

```
[student@serverX ~]$ lab ipv6 grade
IPv4 address is correct ..... PASS
... Output omitted ...
```

Solution

In this lab, you will assign static IPv4 and IPv6 addresses to a network interface.

Resources:

Machines:	serverX
------------------	----------------

Outcomes:

serverX will have a NetworkManager connection, called **eno1**, that manages the **eno1** network interface. It will be configured statically with an IPv4 address of 192.168.0.100/24 and an IPv6 address of fddb:fe2a:ab1e::c0a8:64/64.

Before you begin...

- Reset the serverX system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab ipv6 setup
```

- Become the **root** user.

```
[student@serverX ~]$ sudo -i
```

1. Create a NetworkManager connection, called **eno1**, that corresponds to the **eno1** network interface.

```
[root@serverX ~]# nmcli con add con-name eno1 type ethernet ifname eno1
Connection 'eno1' (0d687259-c64b-4e5b-bece-cabbe952e46f) successfully added.
[root@serverX ~]# nmcli con show
```

NAME	UUID	TYPE	DEVICE
eno1	0d687259-c64b-4e5b-bece-cabbe952e46f	802-3-ethernet	eno1
System eth0	5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03	802-3-ethernet	eth0

2. Configure **eno1** with a static IPv4 address of 192.168.0.100/24.

```
[root@serverX ~]# nmcli con mod eno1 ipv4.addresses 192.168.0.100/24
[root@serverX ~]# nmcli con mod eno1 ipv4.method manual
```

3. Configure **eno1** with a static IPv6 address of fddb:fe2a:ab1e::c0a8:64/64.

```
[root@serverX ~]# nmcli con mod eno1 ipv6.addresses fddb:fe2a:ab1e::c0a8:64/64
[root@serverX ~]# nmcli con mod eno1 ipv6.method manual
```

4. Restart the **eno1** network interface and confirm its new IPv4 and IPv6 address configuration.

- 4.1. Bounce the **eno1** interface by taking it down, then bringing it back up.

```
[root@serverX ~]# nmcli con down eno1
[root@serverX ~]# nmcli con up eno1
Connection successfully activated (D-Bus active path: /org/freedesktop/
NetworkManager/ActiveConnection/2)
```

- 4.2. Use the **ip addr** command to confirm the interface's configuration.

```
[root@serverX ~]# ip addr show dev eno1
4: eno1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
qlen 1000
    link/ether ca:8a:8f:84:e4:8f brd ff:ff:ff:ff:ff:ff
    inet 192.168.0.100/24 brd 192.168.0.255 scope global eno1
        valid_lft forever preferred_lft forever
    inet6 fddb:fe2a:ab1e::c0a8:64/64 scope global
        valid_lft forever preferred_lft forever
    inet6 fe80::c88a:8fff:fe84:e48f/64 scope link
        valid_lft forever preferred_lft forever
```

Notice the new inet address entry.

5. Ping the local IPv4 gateway, 192.168.0.254, and the local IPv6 gateway, fddb:fe2a:ab1e::c0a8:fe, to confirm they can both be reached through the **eno1** interface.

```
[student@serverX ~]$ ping -I eno1 192.168.0.254
PING 192.168.0.254 (192.168.0.254) from 192.168.0.100 eno1: 56(84) bytes of data.
64 bytes from 192.168.0.254: icmp_seq=1 ttl=64 time=0.099 ms
64 bytes from 192.168.0.254: icmp_seq=2 ttl=64 time=0.056 ms
64 bytes from 192.168.0.254: icmp_seq=3 ttl=64 time=0.036 ms
64 bytes from 192.168.0.254: icmp_seq=4 ttl=64 time=0.047 ms
^C
--- 192.168.0.254 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 2999ms
rtt min/avg/max/mdev = 0.036/0.059/0.099/0.025 ms
[student@serverX ~]$ ping6 -I eno1 fddb:fe2a:ab1e::c0a8:fe
PING fddb:fe2a:ab1e::c0a8:fe(fddb:fe2a:ab1e::c0a8:fe) from fddb:fe2a:ab1e::c0a8:64
eno1: 56 data bytes
64 bytes from fddb:fe2a:ab1e::c0a8:fe: icmp_seq=1 ttl=64 time=0.150 ms
64 bytes from fddb:fe2a:ab1e::c0a8:fe: icmp_seq=2 ttl=64 time=0.103 ms
64 bytes from fddb:fe2a:ab1e::c0a8:fe: icmp_seq=3 ttl=64 time=0.049 ms
^C
--- fddb:fe2a:ab1e::c0a8:fe ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2003ms
rtt min/avg/max/mdev = 0.049/0.100/0.150/0.042 ms
```

6. Run the grading script to verify your work.

```
[student@serverX ~]$ lab ipv6 grade
IPv4 address is correct ..... PASS
... Output omitted ...
```

Summary

Review of IPv4 Networking Configuration

In this section, students learned how to:

- Configure IPv4 networking using **nmcli**.
- Configure IPv4 networking using **/etc/sysconfig/network-scripts/ifcfg-*** files.

IPv6 Networking Concepts

In this section, students learned how to:

- Explain the basic details of IPv6 networking.
- Interpret text representations of IPv6 addresses.

IPv6 Networking Configuration

In this section, students learned how to:

- Configure IPv6 networking using **nmcli**.
- Configure IPv6 networking using **/etc/sysconfig/network-scripts/ifcfg-*** files.



CHAPTER 3

CONFIGURING LINK AGGREGATION AND BRIDGING

Overview	
Goal	To configure and troubleshoot advanced network interface functionality, including teaming and local software bridges.
Objectives	<ul style="list-style-type: none">• Use network teaming to provide link redundancy or higher throughput.• Manage a network team interface.• Manage local software bridges and associated interfaces.
Sections	<ul style="list-style-type: none">• Configuring Network Teaming (and Practice)• Managing Network Teaming (and Practice)• Configuring Software Bridges (and Practice)
Lab	<ul style="list-style-type: none">• Configuring Link Aggregation and Bridging

Configuring Network Teaming

Objectives

After completing this section, students should be able to combine two network links using network teaming to provide link redundancy or higher throughput.

Network teaming

Network teaming is method for linking NICs together logically to allow for failover or higher throughput. Teaming is a new implementation that does not affect the older bonding driver in the Linux kernel; it offers an alternate implementation. Red Hat Enterprise Linux 7 supports channel bonding for backward compatibility. Network teaming provides better performance and is more extensible because of its modular design.

Red Hat Enterprise Linux 7 implements network teaming with a small kernel driver and a user-space daemon, **teamd**. The kernel handles network packets efficiently and **teamd** handles logic and interface processing. Software, called *runners*, implement load balancing and active-backup logic, such as **roundrobin**. The following runners are available to **teamd**:

- **broadcast**: a simple runner which transmits each packet from all ports.
- **roundrobin**: a simple runner which transmits packets in a round-robin fashion from each of the ports.
- **activebackup**: this is a failover runner which watches for link changes and selects an active port for data transfers.
- **loadbalance**: this runner monitors traffic and uses a hash function to try to reach a perfect balance when selecting ports for packet transmission.
- **lacp**: implements the 802.3ad Link Aggregation Control Protocol. Can use the same transmit port selection possibilities as the **loadbalance** runner.

All network interaction is done through a *team* interface, composed of multiple *network port* interfaces. When controlling teamed port interfaces using NetworkManager, and especially when fault finding, keep the following in mind:

- Starting the network team interface does not automatically start the port interfaces.
- Starting a port interface always starts the teamed interface.
- Stopping the teamed interface also stops the port interfaces.
- A teamed interface without ports can start static IP connections.
- A team without ports waits for ports when starting DHCP connections.
- A team with a DHCP connection waiting for ports completes when a port with a carrier is added.
- A team with a DHCP connection waiting for ports continues waiting when a port without a carrier is added.

Configuring network teams

The **nmcli** command can be used to create and manage team and port interfaces. The following four steps are used to create and activate a network team interface:

1. Create the team interface.
2. Determine the IPv4 and/or IPv6 attributes of the team interface.
3. Assign the port interfaces.
4. Bring the team and port interfaces up/down.

Create the team interface

Use the **nmcli** command to create a connection for the network team interface, with the following syntax:

```
nmcli con add type team con-name CNAME ifname INAME [config JSON]
```

where *CNAME* will be the name used to refer to the connection, *INAME* will be the interface name, and *JSON* specifies the runner to be used. *JSON* has the following syntax:

```
'{"runner": {"name": "METHOD"}}'
```

where *METHOD* is one of the following: **broadcast**, **roundrobin**, **activebackup**, **loadbalance**, or **lacp**.

Example:

```
[root@demo ~]# nmcli con add type team con-name team0 ifname team0 config '{"runner": {"name": "loadbalance"}}'
```

Determine the IPv4/IPv6 attributes of the team interface

Once the network team interface is created, IPv4 and/or IPv6 attributes can be assigned to it. If DHCP is available, this step is optional, because the default attributes configure the interface to get its IP settings using DHCP.

The following example demonstrates how to assign a static IPv4 address to the **team0** interface:

```
[root@demo ~]# nmcli con mod team0 ipv4.addresses 1.2.3.4/24
[root@demo ~]# nmcli con mod team0 ipv4.method manual
```

Note that the **ipv4.addresses** have to be assigned before the **ipv4.method** can be set to **manual**.

Assign the port interfaces

Use the **nmcli** command to create each of the port interfaces with the following syntax:

```
nmcli con add type team-slave con-name CNAME ifname INAME master TEAM
```

where *CNAME* will be the name used to refer to the port, *INAME* is the name of an existing interface, and *TEAM* specifies the connection name of the network team interface.

The connection name can be explicitly specified, or it will be **team-slave-IFACE** by default.

Example:

```
[root@demo ~]# nmcli con add type team-slave ifname eth1 master team0
[root@demo ~]# nmcli con add type team-slave ifname eth2 master team0 con-name team0-eth2
```

Bring the team and port interfaces up/down

nmcli command can also be used to manage the connections for the team and port interfaces with the following syntax:

```
nmcli dev dis INAME
nmcli con up CNAME
```

INAME is the device name of the team or port interface to be managed. *CNAME* is the connection name of that interface. where *CNAME* is the connection name of the team or port interface to be managed.

Example:

```
[root@demo ~]# nmcli con up team0
[root@demo ~]# nmcli dev dis eth2
```

When the team interface is up, the **teamdctl** command can be used to display the team's state. The output of this command includes the status of the port interfaces.

```
[root@demo ~]# teamdctl team0 state
setup:
  runner: roundrobin
ports:
  eth1
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
  eth2
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
```

R

References

nmcli-examples(5), **teamdctl**(8), and **teamd**(8) man pages

Practice: Configuring Network Teaming

Guided exercise

In this lab, you will create a network team interface.

Resources:	
Machines:	serverX

Outcomes:

A network team interface called **team0** will have a static IP address of 192.168.0.100/24 and will be built upon two port interfaces: **eno1** and **eno2**. It will be a fault-tolerant/active-backup interface.

Before you begin...

- Reset the serverX system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab teambridge setup
```

- Become the **root** user.

```
[student@serverX ~]$ sudo -i
```

- 1. Display the current state of the existing network interfaces. **eno1** and **eno2** will be the interfaces that will be the ports for the teamed interface.

```
[root@serverX ~]# ip link
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
    mode DEFAULT qlen 1000
    link/ether 52:54:00:00:XX:0b brd ff:ff:ff:ff:ff:ff
4: eno1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
    mode DEFAULT qlen 1000
    link/ether 00:10:18:2b:98:85 brd ff:ff:ff:ff:ff:ff
6: eno2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
    mode DEFAULT qlen 1000
    link/ether 64:31:50:18:80:8f brd ff:ff:ff:ff:ff:ff
```

- 2. Create a active-backup teaming interface called **team0** and assign its IPv4 settings.
 - 2.1. Create the **team0** connection.

```
[root@serverX ~]# nmcli con add type team con-name team0 ifname team0
config '{"runner": {"name": "activebackup"}}'
Connection 'team0' (5dc435ac-e4ac-403a-8e8f-163b163bf49b) successfully
added.
```

- 2.2. Define the IPv4 settings for **team0**. Assign it the IP address 192.168.0.100/24; the method for IP should be static.

```
[root@serverX ~]# nmcli con mod team0 ipv4.addresses '192.168.0.100/24'
[root@serverX ~]# nmcli con mod team0 ipv4.method manual
```

- 3. Assign **eno1** and **eno2** as port interfaces for **team0**.

```
[root@serverX ~]# nmcli con add type team-slave con-name team0-port1 ifname eno1
master team0
Connection 'team0-port1' (f5664c4e-1dba-43f8-8427-35aee0594ed3) successfully
added.
[root@serverX ~]# nmcli con add type team-slave con-name team0-port2 ifname eno2
master team0
Connection 'team0-port2' (174e4402-b169-47d1-859f-9a4b3f30000f) successfully
added.
```

- 4. Check the current state of the teamed interfaces on the system.

```
[root@serverX ~]# teamdctl team0 state
setup:
  runner: activebackup
ports:
  eno1
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
  eno2
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
runner:
  active port: eno1
```

Notice how the teamed interface immediately becomes active with **eno1** as the active port.

- 5. Open another terminal and **ping** the local network gateway through the **team0** interface. Let this **ping** continue to run as you perform the following steps.

```
[student@serverX ~]$ ping -I team0 192.168.0.254
PING 192.168.0.254 (192.168.0.254) from 192.168.0.100 team0: 56(84) bytes of
data.
64 bytes from 192.168.0.254: icmp_seq=10 ttl=64 time=1.08 ms
64 bytes from 192.168.0.254: icmp_seq=11 ttl=64 time=0.789 ms
64 bytes from 192.168.0.254: icmp_seq=12 ttl=64 time=0.906 ms
...Output omitted...
```

- 6. Go back to the other **root** terminal. Bring the active port of the teamed interface down and see its impact upon **team0**.

```
[root@serverX ~]# nmcli dev dis eno1
[root@serverX ~]# teamdctl team0 state
setup:
  runner: activebackup
ports:
  eno2
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
runner:
  active port: eno2
```

The **ping** continues to work because **team0** switched over to the remaining port.

- 7. Bring the original port interface back up and bring the other port interface down.

```
[root@serverX ~]# nmcli con up team0-port1
Connection successfully activated (D-Bus active path: /org/freedesktop/
NetworkManager/ActiveConnection/10)
[root@serverX ~]# nmcli dev dis eno2
[root@serverX ~]# teamdctl team0 state
setup:
  runner: activebackup
ports:
  eno1
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
runner:
  active port: eno1
```

Again, note how the **ping** continues to reach the local gateway.

- 8. Bring the down port interface back up and observe how it affects the teamed interface, **team0**.


```
[root@serverX ~]# nmcli con up team0-port2
Connection successfully activated (D-Bus active path: /org/freedesktop/
NetworkManager/ActiveConnection/11)
[root@serverX ~]# teamdctl team0 state
setup:
  runner: activebackup
ports:
  eno1
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
  eno2
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
runner:
  active port: eno1
```

The **ping** continues to contact the local network gateway, and the currently active port interface does not change when both port interfaces are available.

Managing Network Teaming

Objectives

After completing this section, students should be able to troubleshoot and change the configuration of a team interface.

Network teaming configuration files

NetworkManager creates configuration files for network teaming in the `/etc/sysconfig/network-scripts` the same way it does for other interfaces. A configuration file is created for each of the interfaces: for the team, and each of the ports.

The configuration file for the team interface defines the IP settings for the interface. The **DEVICETYPE** variable informs the initialization scripts this is a network team interface. Parameters for **teamd** configuration are defined in the **TEAM_CONFIG** variable. Note that the contents of **TEAM_CONFIG** uses JSON syntax.

```
# /etc/sysconfig/network-scripts/ifcfg-team0
DEVICE=team0
DEVICETYPE=Team
TEAM_CONFIG="{\"runner\": {\"name\": \"broadcast\"}}"
BOOTPROTO=none
IPADDR0=172.25.5.100
PREFIX0=24
NAME=team0
ONBOOT=yes
```

The following is an example configuration file for a team port interface.

```
# /etc/sysconfig/network-scripts/ifcfg-team0-eth1
DEVICE=eth1
DEVICETYPE=TeamPort
TEAM_MASTER=team0
NAME=team0-eth1
ONBOOT=yes
```

The **DEVICETYPE** variable informs the initialization scripts this is a team port interface. The **TEAM_MASTER** variable defines which team device it is a port for.

Setting and adjusting team configuration

Initial network team configuration is set when the team interface is created. The default runner is **roundrobin**, but a different runner can be chosen by specifying a JSON string when the team is created with the **team.config** subcommand. Default values for runner parameters are used when they are not specified.

A different runner can be assigned to an existing team, or runner parameters can be adjusted using the **nmcli con mod** command. The configuration changes can be specified as a JSON string (in the case of simple changes) or the name of a file with a more complex JSON configuration can be given.

```
nmcli con mod IFACE team.config JSON-configuration-file-or-string
```

The following example shows how to assign different priorities to port interfaces in an **active-backup** team:

```
[root@demo ~]# cat /tmp/team.conf
{
  "device": "team0",
  "mcast_rejoin": {
    "count": 1
  },
  "notify_peers": {
    "count": 1
  },
  "ports": {
    "eth1": {
      "prio": -10,
      "sticky": true,
      "link_watch": {
        "name": "ethtool"
      }
    },
    "eth2": {
      "prio": 100,
      "link_watch": {
        "name": "ethtool"
      }
    }
  },
  "runner": {
    "name": "activebackup"
  }
}
[root@demo ~]# nmcli con mod team0 team.config /tmp/team.conf
```



Note

Any changes made do not go into effect until the next time the team interface is brought up.

The **link_watch** settings in the configuration file determines how the link state of the port interfaces are monitored. The default looks like the following, and uses functionality similar to the **ethtool** command to check the link of each interface:

```
"link_watch": {
  "name": "ethtool"
}
```

Another way to check link state is to periodically use an ARP ping packet to check for remote connectivity. Local and remote IP addresses and timeouts would have to be specified. A configuration that would accomplish that would look similar to the following:

```
"link_watch":{
  "name": "arp_ping",
  "interval": 100,
  "missed_max": 30,
  "source_host": "192.168.23.2",
  "target_host": "192.168.23.1"
},
```



Note

Be aware that omitted options revert to their default values when they are not specified in the JSON file.

Troubleshooting network teams

The **teamnl** and **teamdctl** commands are very useful for troubleshooting network teams. These commands only work on network teams that are up. The following examples show some typical uses for these commands.

Display the team ports of the **team0** interface:

```
[root@demo ~]# teamnl team0 ports
4: eth2: up 0Mbit HD
3: eth1: up 0Mbit HD
```

Display the currently active port of **team0**:

```
[root@demo ~]# teamnl team0 getoption activeport
3
```

Set the option for the active port of **team0**:

```
[root@demo ~]# teamnl team0 setoption activeport 3
```

Use **teamdctl** to display the current state of the **team0** interface:

```
[root@demo ~]# teamdctl team0 state
setup:
  runner: activebackup
ports:
  eth2
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
  eth1
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
runner:
  active port: eth1
```

Use **teamdctl** to display the current JSON configuration for **team0**:

```
[root@demo ~]# teamdctl team0 config dump
{
  "device": "team0",
  "mcast_rejoin": {
    "count": 1
  },
  "notify_peers": {
    "count": 1
  },
  "ports": {
    "eth1": {
      "link_watch": {
        "name": "ethtool"
      },
      "prio": -10,
      "sticky": true
    },
    "eth2": {
      "link_watch": {
        "name": "ethtool"
      },
      "prio": 100
    }
  },
  "runner": {
    "name": "activebackup"
  }
}
```

R

References

teamd.conf(5) and **teamnl**(8) man pages

Practice: Managing Network Teaming

Guided exercise

In this lab, you will manage your network team interface. You will deactivate it, change its runner to **roundrobin**, and reactivate it. You will use **teamdctl** and **teamnl** to get information about the team interface.

Resources:	
Files:	<code>/etc/sysconfig/network-scripts/ifcfg-*</code>
Machines:	<code>serverX</code>

Outcomes:

A network team interface called **team0** will have a static IP address of 192.168.0.100/24 and will be built upon two port interfaces: **eno1** and **eno2**. Its runner will be changed to **roundrobin**, then changed back to **activebackup**.

Before you begin...

- Complete the Configuring Network Teaming practice exercise in the previous section. This exercise uses the teamed interfaces created in that practice exercise.
- Log into your server system and become **root**.

```
[student@serverX ~]$ sudo -i
```

- 1. Get the initial state of the teamed interfaces on the system.

```
[root@serverX ~]# teamdctl team0 state
setup:
  runner: activebackup
ports:
  eno1
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
  eno2
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
runner:
  active port: eno1
```

- 2. Examine the configuration files created by NetworkManager that apply to the team interfaces and its ports.
 - 2.1. Display the file for the team interface and note how it defines the runner to be used and the IPv4 network settings for the interface.

```
[root@serverX ~]# cat /etc/sysconfig/network-scripts/ifcfg-team0
DEVICE=team0
TEAM_CONFIG="{\"runner\": {\"name\": \"activebackup\"}}"
DEVICETYPE=Team
BOOTPROTO=none
DEFROUTE=yes
IPV4_FAILURE_FATAL=no
IPV6INIT=yes
IPV6_AUTOCONF=yes
IPV6_DEFROUTE=yes
IPV6_FAILURE_FATAL=no
NAME=team0
UUID=5dc435ac-e4ac-403a-8e8f-163b163bf49b
ONBOOT=yes
IPADDR0=192.168.0.100
PREFIX0=24
IPV6_PEERDNS=yes
IPV6_PEERROUTES=yes
```

- 2.2. Display the configuration files for the port interfaces. Take special notice of the values of the **TEAM_MASTER** and **DEVICETYPE** shell variables.

```
[root@serverX ~]# cat /etc/sysconfig/network-scripts/ifcfg-team0-port1
BOOTPROTO=dhcp
DEFROUTE=yes
PEERDNS=yes
PEERROUTES=yes
IPV4_FAILURE_FATAL=no
IPV6INIT=yes
IPV6_AUTOCONF=yes
IPV6_DEFROUTE=yes
IPV6_PEERDNS=yes
IPV6_PEERROUTES=yes
IPV6_FAILURE_FATAL=no
NAME=team0-port1
UUID=f5664c4e-1dba-43f8-8427-35aee0594ed3
DEVICE=eno1
ONBOOT=yes
TEAM_MASTER=team0
DEVICETYPE=TeamPort
[root@serverX ~]# cat /etc/sysconfig/network-scripts/ifcfg-team0-port2
BOOTPROTO=dhcp
DEFROUTE=yes
PEERDNS=yes
PEERROUTES=yes
IPV4_FAILURE_FATAL=no
IPV6INIT=yes
IPV6_AUTOCONF=yes
IPV6_DEFROUTE=yes
IPV6_PEERDNS=yes
IPV6_PEERROUTES=yes
IPV6_FAILURE_FATAL=no
NAME=team0-port2
UUID=174e4402-b169-47d1-859f-9a4b3f30000f
DEVICE=eno2
ONBOOT=yes
TEAM_MASTER=team0
DEVICETYPE=TeamPort
```

- 3. Bring the **team0** interface down and edit the configuration file to use the **roundrobin** runner.

- 3.1. Bring the **team0** interface down.

```
[root@serverX ~]# nmcli dev dis team0
```

- 3.2. Edit the configuration file for **team0** and adjust it to use the **roundrobin** runner.

```
[root@serverX ~]# vim /etc/sysconfig/network-scripts/ifcfg-team0
[root@serverX ~]# grep runner /etc/sysconfig/network-scripts/ifcfg-team0
TEAM_CONFIG="{\"runner\": {\"name\": \"roundrobin\"}}"
```

- 3.3. Use **nmcli** to make NetworkManager reload the updated configuration.

```
[root@serverX ~]# nmcli con load /etc/sysconfig/network-scripts/ifcfg-team0
```

- 4. Bring the **team0** interface back up.

- 4.1. Tell NetworkManager to activate **team0**.

```
[root@serverX ~]# nmcli con up team0
Connection successfully activated (D-Bus active path: /org/freedesktop/NetworkManager/ActiveConnection/12)
```

- 4.2. Check the current state of **team0**.

```
[root@serverX ~]# teamdctl team0 state
setup:
runner: roundrobin
```

- 4.3. Log into another window and **ping** the local network gateway through the **team0** interface.

```
[root@serverX ~]# ping -I team0 192.168.0.254
PING 192.168.0.254 (192.168.0.254) from 192.168.0.100 team0: 56(84) bytes of data.
From 192.168.0.100 icmp_seq=1 Destination Host Unreachable
From 192.168.0.100 icmp_seq=2 Destination Host Unreachable
From 192.168.0.100 icmp_seq=3 Destination Host Unreachable
...Output omitted...
```

The **ping** command fails to contact the gateway because the teamed interface doesn't have any active ports.

- 5. Activate one of the port interfaces for **team0**.

- 5.1. Use **nmcli** to activate the **team0-port1** interface.


```
[root@serverX ~]# nmcli con up team0-port1
Connection successfully activated (D-Bus active path: /org/freedesktop/
NetworkManager/ActiveConnection/13)
```

- 5.2. Display the state of the teamed interface.

```
[root@serverX ~]# teamdctl team0 state
setup:
  runner: roundrobin
ports:
  eno1
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
```

- 5.3. **ping** the local network gateway through the **team0** interface. It should be able to reach the gateway.

```
[root@serverX ~]# ping -I team0 192.168.0.254
PING 192.168.0.254 (192.168.0.254) from 192.168.0.100 team0: 56(84) bytes
of data.
64 bytes from 192.168.0.254: icmp_seq=1 ttl=64 time=0.516 ms
64 bytes from 192.168.0.254: icmp_seq=2 ttl=64 time=0.703 ms
64 bytes from 192.168.0.254: icmp_seq=3 ttl=64 time=0.422 ms
```

- 6. Bring up the other team port for **team0**.

- 6.1. Use **nmcli** to activate the **team0-port2** interface.

```
[root@serverX ~]# nmcli con up team0-port2
Connection successfully activated (D-Bus active path: /org/freedesktop/
NetworkManager/ActiveConnection/14)
```

- 6.2. Display the updated state of the teamed interface.

```
[root@serverX ~]# teamdctl team0 state
setup:
  runner: roundrobin
ports:
  eno1
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
  eno2
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
```

- 7. Use **teamdctl** to display the configuration for **team0**.

```
[root@serverX ~]# teamdctl team0 config dump
{
  "device": "team0",
  "ports": {
    "eno1": {
      "link_watch": {
        "name": "ethtool"
      }
    },
    "eno2": {
      "link_watch": {
        "name": "ethtool"
      }
    }
  },
  "runner": {
    "name": "roundrobin"
  }
}
```

- 8. Use the **teamnl** command to display the tunable options for **team0**.

```
[root@serverX ~]# teamnl team0 options
queue_id (port:eno2) 0
priority (port:eno2) 0
user_linkup_enabled (port:eno2) false
user_linkup (port:eno2) true
enabled (port:eno2) true
queue_id (port:eno1) 0
priority (port:eno1) 0
user_linkup_enabled (port:eno1) false
user_linkup (port:eno1) true
enabled (port:eno1) true
mcast_rejoin_interval 0
mcast_rejoin_count 0
notify_peers_interval 0
notify_peers_count 0
mode roundrobin
```

- 9. Modify the **team0** interface so it uses the **activebackup** runner instead of **roundrobin**.

- 9.1. The interface can only be modified after it is brought down.

```
[root@serverX ~]# nmcli dev dis team0
```

- 9.2. Use **nmcli** to tune the teamed interface to use the **activebackup** runner.

```
[root@serverX ~]# nmcli con mod team0 team.config '{"runner": {"name": "activebackup"}}'
```

- 9.3. Examine the changes made to the interface's configuration file.

```
[root@serverX ~]# cat /etc/sysconfig/network-scripts/ifcfg-team0
DEVICE=team0
TEAM_CONFIG="{\"runner\": {\"name\": \"activebackup\"}}"
DEVICETYPE=Team
BOOTPROTO=none
DEFROUTE=yes
IPV4_FAILURE_FATAL=no
IPV6INIT=yes
IPV6_AUTOCONF=yes
IPV6_DEFROUTE=yes
IPV6_FAILURE_FATAL=no
NAME=team0
UUID=5dc435ac-e4ac-403a-8e8f-163b163bf49b
ONBOOT=yes
IPADDR0=192.168.0.100
PREFIX0=24
IPV6_PEERDNS=yes
IPV6_PEERROUTES=yes
```

- 10. Reactivate the teamed interface and both of its port interfaces.
 - 10.1. Use **nmcli** to activate the teamed interface.

```
[root@serverX ~]# nmcli con up team0
Connection successfully activated (D-Bus active path: /org/freedesktop/NetworkManager/ActiveConnection/15)
```

- 10.2. Display its initial state. Note that the port interfaces didn't get activated.

```
[root@serverX ~]# teamdctl team0 state
setup:
  runner: activebackup
runner:
  active port:
```

- 10.3. Activate the two port interfaces and display the resulting teamed interface state.

```
[root@serverX ~]# nmcli con up team0-port1
Connection successfully activated (D-Bus active path: /org/freedesktop/
NetworkManager/ActiveConnection/16)
[root@serverX ~]# nmcli con up team0-port2
Connection successfully activated (D-Bus active path: /org/freedesktop/
NetworkManager/ActiveConnection/17)
[root@serverX ~]# teamdctl team0 state
setup:
  runner: activebackup
ports:
  eno1
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
  eno2
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
runner:
  active port: eno1
```

Restart **ping** and make sure it works once the first port interface is brought back up.

```
[root@serverX ~]# ping -I team0 192.168.0.254
PING 192.168.0.254 (192.168.0.254) from 192.168.0.100 team0: 56(84) bytes
of data.
64 bytes from 192.168.0.254: icmp_seq=1 ttl=64 time=0.516 ms
64 bytes from 192.168.0.254: icmp_seq=2 ttl=64 time=0.703 ms
64 bytes from 192.168.0.254: icmp_seq=3 ttl=64 time=0.422 ms
```

- 11. Use the **teamnl** to display the options available to an **activebackup** team device.

```
[root@serverX ~]# teamnl team0 options
queue_id (port:eno2) 0
priority (port:eno2) 0
user_linkup_enabled (port:eno2) false
user_linkup (port:eno2) true
enabled (port:eno2) false
queue_id (port:eno1) 0
priority (port:eno1) 0
user_linkup_enabled (port:eno1) false
user_linkup (port:eno1) true
enabled (port:eno1) true
activeport 3
mcast_rejoin_interval 0
mcast_rejoin_count 1
notify_peers_interval 0
notify_peers_count 1
mode activebackup
```

Configuring Software Bridges

Objectives

After completing this section, students should be able to configure and troubleshoot local software bridges and associated interfaces.

Software bridges

A network bridge is a link-layer device that forwards traffic between networks based on MAC addresses. It learns what hosts are connected to each network, builds a table of MAC addresses, then makes packet forwarding decisions based on that table. A software bridge can be used in a Linux environment to emulate a hardware bridge. The most common application for software bridges is in virtualization applications for sharing a hardware NIC among one or more virtual NICs.

Configure software bridges

The **nmcli** can be used to configure software bridges persistently. First, the software bridge is created, then existing interfaces are connected to it. For example, the following commands will create a bridge called **br0** and attach both the **eth1** and **eth2** interfaces to it.

```
[root@demo ~]# nmcli con add type bridge con-name br0 ifname br0
[root@demo ~]# nmcli con add type bridge-slave con-name br0-port1 ifname eth1 master br0
[root@demo ~]# nmcli con add type bridge-slave con-name br0-port2 ifname eth2 master br0
```



Note

NetworkManager can only attach Ethernet interfaces to a bridge. It does not support aggregate interfaces, such as a teamed or bonded interface. These must be configured by manipulating the configuration files in **/etc/sysconfig/network-scripts**.

Software bridge configuration files

Software bridges are managed by interface configuration files found in the **/etc/sysconfig/network-scripts** directory. There is an **ifcfg-*** configuration file for each software bridge.

The following is a sample configuration file for a software bridge:

```
# /etc/sysconfig/network-scripts/ifcfg-br1
DEVICE=br1
NAME=br1
TYPE=Bridge
BOOTPROTO=none
IPADDR0=172.25.5.100
PREFIX0=24
STP=yes
BRIDGING_OPTS=priority=32768
```

The **TYPE=Bridge** definition specifies that this is a software bridge. **BRIDGING_OPTS** defines additional bridge options. Note that this bridge has been assigned a static IP address.

The following configuration file attaches an Ethernet interface to a software bridge:

```
# /etc/sysconfig/network-scripts/ifcfg-br1-port0
TYPE=Ethernet
NAME=br1-port0
DEVICE=eth1
ONBOOT=yes
BRIDGE=br1
```

The single variable definition, **BRIDGE=br1**, is what ties this interface to the software bridge, **br1**.



Note

To implement a software bridge on an existing teamed or bonded network interface managed by NetworkManager, NetworkManager will have to be disabled since it only supports bridges on simple Ethernet interfaces. Configuration files for the bridge will have to be created by hand. The **ifup** and **ifdown** utilities can be used to manage the software bridge and other network interfaces.

The **brctl show** command will display software bridges and the list of interfaces attached to them.

```
[root@demo ~]# brctl show
bridge name      bridge id                STP enabled  interfaces
br1              8000.52540001050b       yes          eth1
```

R

References

nmcli-examples(5) and **brctl(8)** man pages

Practice: Configuring Software Bridges

Guided exercise

In this lab, you will create a network bridge.

Resources:	
Files:	<code>/etc/sysconfig/network-scripts/ifcfg-*</code>
Machines:	<code>serverX</code>

Outcomes:

A network bridge called **br1** will be attached to the **eno1** network interface. It will have a static IP address of 192.168.0.100/24.

Before you begin...

- Reset the serverX system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab teambridge setup
```

- Become the **root** user.

```
[student@serverX ~]$ sudo -i
```

- 1. Define a software bridge called **br1** and assign it a static IPv4 address of 192.168.0.100/24.

- 1.1. Use **nmcli** to create the software bridge.

```
[root@serverX ~]# nmcli con add type bridge con-name br1 ifname br1
Connection 'br1' (d9d56520-574a-4e2a-9f43-b593a1bdff61) successfully
added.
```

- 1.2. Configure the IPv4 addressing for the interface.

```
[root@serverX ~]# nmcli con mod br1 ipv4.addresses 192.168.0.100/24
[root@serverX ~]# nmcli con mod br1 ipv4.method manual
```

- 2. Attach the **eno1** interface to the **br1** software bridge.

```
[root@serverX ~]# nmcli con add type bridge-slave con-name br1-port0 ifname eno1
master br1
Connection 'br1-port0' (5f5e7ea8-b507-4c10-a61f-779369cf82ee) successfully added.
```

- 3. Inspect the configuration files that were created for the software bridge by NetworkManager. Look for variables that connect the two interfaces.

```
[root@serverX ~]# cat /etc/sysconfig/network-scripts/ifcfg-br1
DEVICE=br1
STP=yes
TYPE=Bridge
BOOTPROTO=none
DEFROUTE=yes
IPV4_FAILURE_FATAL=no
IPV6INIT=yes
IPV6_AUTOCONF=yes
IPV6_DEFROUTE=yes
IPV6_FAILURE_FATAL=no
NAME=br1
UUID=d9d56520-574a-4e2a-9f43-b593a1bdff61
ONBOOT=yes
IPADDR0=192.168.0.100
PREFIX0=24
BRIDGING_OPTS=priority=32768
IPV6_PEERDNS=yes
IPV6_PEERROUTES=yes
[root@serverX ~]# cat /etc/sysconfig/network-scripts/ifcfg-br1-port0
TYPE=Ethernet
NAME=br1-port0
UUID=5f5e7ea8-b507-4c10-a61f-779369cf82ee
DEVICE=eno1
ONBOOT=yes
BRIDGE=br1
```

- 4. Display the link status of the network interfaces.

```
[root@serverX ~]# ip link
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
    mode DEFAULT qlen 1000
    link/ether 52:54:00:00:XX:0b brd ff:ff:ff:ff:ff:ff
4: eno1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast master br1
    state UP mode DEFAULT qlen 1000
    link/ether 00:10:18:2b:98:85 brd ff:ff:ff:ff:ff:ff
6: eno2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
    mode DEFAULT qlen 1000
    link/ether 64:31:50:18:80:8f brd ff:ff:ff:ff:ff:ff
7: br1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP mode
    DEFAULT
    link/ether 00:10:18:2b:98:85 brd ff:ff:ff:ff:ff:ff
```

- 5. Use **brctl** to display information about software bridges on the system.

```
[root@serverX ~]# brctl show
bridge name      bridge id                STP enabled  interfaces
br1              8000.52540001050b       yes          eno1
```

- 6. Ping the local network gateway using the software bridge.


```
[root@serverX ~]# ping -I br1 192.168.0.254
PING 192.168.0.254 (192.168.0.254) from 192.168.0.100 br1: 56(84) bytes of data.
64 bytes from 192.168.0.254: icmp_seq=10 ttl=64 time=0.520 ms
64 bytes from 192.168.0.254: icmp_seq=11 ttl=64 time=0.470 ms
64 bytes from 192.168.0.254: icmp_seq=12 ttl=64 time=0.339 ms
64 bytes from 192.168.0.254: icmp_seq=13 ttl=64 time=0.294 ms
...Output omitted...
```

Lab: Configuring Link Aggregation and Bridging

Performance checklist

In this lab, you will create a bridge that is connected to a network team interface.

Resources:	
Files:	<code>/etc/sysconfig/network-scripts/ifcfg-*</code>
Machines:	<code>serverX</code>

Outcomes:

serverX will have an **activebackup** team interface, called **team0**. The **team0** interface will be built upon the port interfaces **eno1** and **eno2**. **team0** will be attached to a bridge, called **brteam0**. The bridge will have a static IP address of 192.168.0.100/24.

Before you begin...

- Reset the serverX system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab teambridge setup
```

- Become the **root** user.

```
[student@serverX ~]$ sudo -i
```

1. Confirm that **eno1** and **eno2** are available for use. Display the current state of the existing network interfaces. **eno1** and **eno2** will be the port interfaces that will be teamed into a single interface.
2. Create an **activebackup** network team interface called **team0**.
3. Disable the **team0** device then the NetworkManager service.
4. Manipulate the interface configuration files so that the **team0** interface is attached to a software bridge called **brteam0**. The bridge should have a static IP address of 192.168.0.100/24.
5. Reset the network to start the new bridge, **brteam0**, and reactivate the **team0** interface.
6. Reboot **serverX**, log in as student, and test the network connectivity of the new bridge, **brteam0**. Verify your configuration by running **lab teambridge grade**.

Solution

In this lab, you will create a bridge that is connected to a network team interface.

Resources:	
Files:	<code>/etc/sysconfig/network-scripts/ifcfg-*</code>
Machines:	<code>serverX</code>

Outcomes:

serverX will have an **activebackup** network team interface, called **team0**. The **team0** interface will be built upon the port interfaces **eno1** and **eno2**. **team0** will be attached to a bridge, called **brteam0**. The bridge will have a static IP address of 192.168.0.100/24.

Before you begin...

- Reset the serverX system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab teambridge setup
```

- Become the **root** user.

```
[student@serverX ~]$ sudo -i
```

1. Confirm that **eno1** and **eno2** are available for use. Display the current state of the existing network interfaces. **eno1** and **eno2** will be the port interfaces that will be teamed into a single interface.

```
[root@serverX ~]# ip link
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP mode
    DEFAULT qlen 1000
    link/ether 52:54:00:00:XX:0b brd ff:ff:ff:ff:ff:ff
4: eno1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP mode
    DEFAULT qlen 1000
    link/ether 00:10:18:2b:98:85 brd ff:ff:ff:ff:ff:ff
6: eno2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP mode
    DEFAULT qlen 1000
    link/ether 64:31:50:18:80:8f brd ff:ff:ff:ff:ff:ff
```

2. Create an **activebackup** network team interface called **team0**.

- 2.1. Create the network team interface and call it **team0**.

```
[root@serverX ~]# nmcli con add type team con-name team0 ifname team0 \
> config '{"runner": {"name": "activebackup"}}'
Connection 'team0' (2f608473-ff8b-4a0d-b250-79567e3f4a13) successfully added.
```

- 2.2. Assign **eno1** and **eno2** as network port interfaces for **team0**.

```
[root@serverX ~]# nmcli con add type team-slave con-name team0-port1 ifname eno1
master team0
Connection 'team0-port1' (3367d0ef-deb5-444b-bc01-0ed3825615a9) successfully
added.
[root@serverX ~]# nmcli con add type team-slave con-name team0-port2 ifname eno2
master team0
Connection 'team0-port2' (4951d25b-a454-4735-8c7b-a51b983df56b) successfully
added.
```

2.3. Confirm the **team0** interface is up and working properly.

```
[root@serverX ~]# teamdctl team0 state
setup:
  runner: activebackup
ports:
  eno1
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
  eno2
  link watches:
    link summary: up
    instance[link_watch_0]:
      name: ethtool
      link: up
runner:
  active port: eno1
```

3. Disable the **team0** device then the NetworkManager service.

```
[root@serverX ~]# nmcli dev dis team0
[root@serverX ~]# systemctl stop NetworkManager
[root@serverX ~]# systemctl disable NetworkManager
rm '/etc/systemd/system/multi-user.target.wants/NetworkManager.service'
rm '/etc/systemd/system/dbus-org.freedesktop.NetworkManager.service'
rm '/etc/systemd/system/dbus-org.freedesktop.nm-dispatcher.service'
```

4. Manipulate the interface configuration files so that the **team0** interface is attached to a software bridge called **brteam0**. The bridge should have a static IP address of 192.168.0.100/24.

4.1. List the original interface configuration files that were created by NetworkManager:

```
[root@serverX ~]# cd /etc/sysconfig/network-scripts/
[root@serverX network-scripts]# ls -1 ifcfg-*
ifcfg-eth0
ifcfg-lo
ifcfg-team0
ifcfg-team0-port1
ifcfg-team0-port2
```

4.2. Edit **ifcfg-team0** and define a **BRIDGE** variable that assigns it to the new bridge about to be created, **brteam0**.

```
[root@serverX network-scripts]# vim ifcfg-team0
[root@serverX network-scripts]# tail ifcfg-team0
IPV6INIT=yes
IPV6_AUTOCONF=yes
IPV6_DEFROUTE=yes
IPV6_PEERDNS=yes
IPV6_PEERROUTES=yes
IPV6_FAILURE_FATAL=no
NAME=team0
UUID=2f608473-ff8b-4a0d-b250-79567e3f4a13
ONBOOT=yes
BRIDGE=br team0
BRIDGE=brteam0
```

- 4.3. Remove the IP configuration information from the configuration files for the team port interfaces.

```
[root@serverX network-scripts]# vim ifcfg-team0-port1
[root@serverX network-scripts]# cat ifcfg-team0-port1
NAME=team0-port1
UUID=3367d0ef-deb5-444b-bc01-0ed3825615a9
DEVICE=eno1
ONBOOT=yes
TEAM_MASTER=team0
DEVICETYPE=TeamPort
```

Repeat these steps for the **ifcfg-team0-port2** configuration file.

- 4.4. Create a new interface configuration file for the bridge, **ifcfg-brteam0**. Define the IP configuration information in that file.

```
[root@serverX network-scripts]# vim ifcfg-brteam0
[root@serverX network-scripts]# cat ifcfg-brteam0
DEVICE=brteam0
ONBOOT=yes
TYPE=Bridge
IPADDR0=192.168.0.100
PREFIX0=24
```

5. Reset the network to start the new bridge, **brteam0**, and reactivate the **team0** interface.

```
[root@serverX network-scripts]# systemctl restart network
```

6. Reboot **serverX**, log in as student, and test the network connectivity of the new bridge, **brteam0**. Verify your configuration by running **lab teambridge grade**.

- 6.1. Reboot **serverX**.

```
[root@serverX network-scripts]# reboot
```

- 6.2. Use **ping** to see if the local gateway, 192.168.0.254, can be reached through the **brteam0** interface.

```
[student@serverX ~]$ ping -I brteam0 192.168.0.254
PING 192.168.0.254 (192.168.0.254) from 192.168.0.100 brteam0: 56(84) bytes of
data.
64 bytes from 192.168.0.254: icmp_seq=1 ttl=64 time=0.172 ms
64 bytes from 192.168.0.254: icmp_seq=2 ttl=64 time=0.091 ms
64 bytes from 192.168.0.254: icmp_seq=3 ttl=64 time=0.052 ms
... Output omitted ...
```

6.3. Run the grading script to verify your work.

```
[student@serverX ~]$ lab teambridge grade
Confirming teamed interface exists .... PASS
It uses specified port interfaces ..... PASS
Correct team runner implemented ..... PASS
... Output omitted ...
```

Summary

Configuring Network Teaming

In this section, students learned how to define network teaming interfaces.

Managing Network Teaming

In this section, students learned how to:

- Navigate and modify network teaming configuration files.
- Modify team runner configuration settings.
- Use **teamnl** and **teamdctl** to troubleshoot network teaming.

Configuring Software Bridges

In this section, students learned how to define software bridge network interfaces.



CHAPTER 4

NETWORK PORT SECURITY

Overview	
Goal	To permit and reject access to network services using advanced SELinux and firewalld filtering techniques.
Objectives	<ul style="list-style-type: none">• Review firewalld concepts and management commands covered in previous courses.• Configure more complex firewall configurations using firewalld's support for "rich language rules."• Describe and implement Network Address Translation (NAT).• Ensure network ports have the correct SELinux type so that services are able to bind to them.
Sections	<ul style="list-style-type: none">• Managing Firewalld (and Practice)• Managing Rich Rules (and Practice)• Masquerading and Port Forwarding (and Practice)• Managing SELinux Port Labeling (and Practice)
Lab	<ul style="list-style-type: none">• Network Port Security

Managing FirewallD

Objectives

After completing this section students should be able to review firewallD concepts and management commands covered in previous courses.

Firewalld overview

firewalld is the default method in Red Hat Enterprise Linux 7 for managing host-level firewalls. Started from the **firewalld.service** **systemd** service, **firewalld** manages the Linux kernel **netfilter** subsystem using the low-level **iptables**, **ip6tables**, and **ebtables** commands.



Note

The **firewalld.service** and **iptables.service**, **ip6tables.service**, and **ebtables.service** services conflict with each other. To prevent accidentally starting one of the ***tables.service** services (and wiping out any running firewall config in the process), it is good practice to *mask* them using **systemctl**.

```
[root@serverX ~]# for SERVICE in iptables ip6tables ebtables; do
> systemctl mask ${SERVICE}.service
> done
```

firewalld separates all incoming traffic into zones, with each zone having its own set of rules. To check which zone to use for an incoming connection, **firewalld** uses this logic, where the first rule that matches wins:

1. If the *source address* of an incoming packet matches a source rule setup for a zone, that packet will be routed through that zone.
2. If the *incoming interface* for a packet matches a filter setup for a zone, that zone will be used.
3. Otherwise, the *default zone* is used. The default zone is not a separate zone; instead, it points to one of the other zones defined on the system.

Unless overridden by an administrator or a **NetworkManager** configuration, the default zone for any new network interface will be set to the **public** zone.

A number of predefined zones are shipped with **firewalld**, each with their own intended usage:

Default Configuration of firewallD Zones

Zone name	Default configuration
trusted	Allow all incoming traffic.
home	Reject incoming traffic unless related to outgoing traffic or matching the ssh , mdns , ipp-client , samba-client , or dhcpv6-client predefined services.

Zone name	Default configuration
internal	Reject incoming traffic unless related to outgoing traffic or matching the ssh , mdns , ipp-client , samba-client , or dhcpv6-client predefined services (same as the home zone to start with).
work	Reject incoming traffic unless related to outgoing traffic or matching the ssh , ipp-client , or dhcpv6-client predefined services.
public	Reject incoming traffic unless related to outgoing traffic or matching the ssh or dhcpv6-client predefined services. <i>The default zone for newly added network interfaces.</i>
external	Reject incoming traffic unless related to outgoing traffic or matching the ssh predefined service. Outgoing IPv4 traffic forwarded through this zone is <i>masqueraded</i> to look like it originated from the IPv4 address of the outgoing network interface.
dmz	Reject incoming traffic unless related to outgoing traffic or matching the ssh predefined service.
block	Reject all incoming traffic unless related to outgoing traffic.
drop	Drop all incoming traffic unless related to outgoing traffic (do not even respond with ICMP errors).

Managing firewalld

firewalld can be managed in three ways:

1. Using the command-line tool **firewall-cmd**.
2. Using the graphical tool **firewall-config**.
3. Using the configuration files in **/etc/firewalld/**.

In most cases, editing the configuration files directly is not recommended, but it can be useful to copy configurations in this way when using configuration management tools.

Configure firewall settings with firewall-cmd

This section will focus on managing **firewalld** using the command-line tool **firewall-cmd**.

firewall-cmd is installed as part of the main *firewalld* package. **firewall-cmd** can perform the same actions as **firewall-config**.

The following table lists a number of frequently used **firewall-cmd** commands, along with an explanation. Note that unless otherwise specified, almost all commands will work on the *runtime* configuration, unless the **--permanent** option is specified. Many of the commands listed take the **--zone=<ZONE>** option to determine which zone they affect. If **--zone** is omitted from those commands, the *default* zone is used.

While configuring a firewall, an administrator will normally apply all changes to the **--permanent** configuration, and then activate those changes with **firewall-cmd --reload**. While testing out new, and possibly dangerous, rules, an administrator can choose to work on the *runtime* configuration by omitting the **--permanent** option. In those cases, an extra option can be used to automatically remove a rule after a certain amount of time, preventing an administrator from accidentally locking out a system: **--timeout=<TIMEINSECONDS>**.

firewall-cmd commands	Explanation
--get-default-zone	Query the current default zone.
--set-default-zone=<ZONE>	Set the default zone. This changes both the runtime and the permanent configuration.
--get-zones	List all available zones.
--get-services	List all predefined services.
--get-active-zones	List all zones currently in use (have an interface or source tied to them), along with their interface and source information.
--add-source=<CIDR> [--zone=<ZONE>]	Route all traffic coming from the IP address or network/netmask <CIDR> to the specified zone. If no --zone= option is provided, the default zone will be used.
--remove-source=<CIDR> [--zone=<ZONE>]	Remove the rule routing all traffic coming from the IP address or network/netmask <CIDR> from the specified zone. If no --zone= option is provided, the default zone will be used.
--add-interface=<INTERFACE> [--zone=<ZONE>]	Route all traffic coming from <INTERFACE> to the specified zone. If no --zone= option is provided, the default zone will be used.
--change-interface=<INTERFACE> [--zone=<ZONE>]	Associate the interface with <ZONE> instead of its current zone. If no --zone= option is provided, the default zone will be used.
--list-all [--zone=<ZONE>]	List all configured interfaces, sources, services, and ports for <ZONE> . If no --zone= option is provided, the default zone will be used.
--list-all-zones	Retrieve all information for all zones (interfaces, sources, ports, services, etc.).
--add-service=<SERVICE>	Allow traffic to <SERVICE> . If no --zone= option is provided, the default zone will be used.
--add-port=<PORT/PROTOCOL>	Allow traffic to the <PORT/PROTOCOL> port(s). If no --zone= option is provided, the default zone will be used.

firewall-cmd commands	Explanation
<code>--remove-service=<SERVICE></code>	Remove <SERVICE> from the allowed list for the zone. If no <code>--zone=</code> option is provided, the default zone will be used.
<code>--remove-port=<PORT/PROTOCOL></code>	Remove the <PORT/PROTOCOL> port(s) from the allowed list for the zone. If no <code>--zone=</code> option is provided, the default zone will be used.
<code>--reload</code>	Drop the runtime configuration and apply the persistent configuration.

firewall-cmd example

The following examples show the default zone being set to **dmz**, all traffic coming from the **192.168.0.0/24** network being assigned to the **internal** zone, and the network ports for **mysql** being opened on the **internal** zone.

```
[root@serverX ~]# firewall-cmd --set-default-zone=dmz
[root@serverX ~]# firewall-cmd --permanent --zone=internal --add-source=192.168.0.0/24
[root@serverX ~]# firewall-cmd --permanent --zone=internal --add-service=mysql
[root@serverX ~]# firewall-cmd --reload
```

Firewalld configuration files

firewalld configuration files are stored in two places: **/etc/firewalld** and **/usr/lib/firewalld**. If a configuration file with the same name is stored in both locations, the version from **/etc/firewalld/** will be used. This allows administrators to override default zones and settings without fear of their changes being wiped out by a package update.

R

References

firewall-cmd(1), **firewall-config(1)**, **firewalld(1)**, **firewalld.zone(5)**, and **firewalld.zones(5)** man pages

Practice: Configuring a Firewall

Guided exercise

In this lab, you will configure a basic `firewalld` firewall.

Resources:	
Machines:	desktopX and serverX

Outcomes:

A firewall configured on **serverX** to allow **https** traffic from your **desktopX** system.

Before you begin...

- Reset your **desktopX** system.
- Reset your **serverX** system.

As part of an ongoing project to track the designated caffeinated beverage fetcher of the day, you have been tasked with configuring a basic web server and firewall on your **serverX** system.

Your setup must meet these requirements:

- The `httpd` and `mod_ssl` packages must be installed on **serverX**.
- The `httpd.service` must be enabled and started on **serverX**.
- Until your in-house web developers complete the actual application, the web server should serve a placeholder page with the text **COFFEE!**.
- `firewalld` must be enabled and started on **serverX**.
- The `firewalld` configuration on **serverX** must use the `dmz` zone for all unspecified connections.
- Traffic from the `172.25.X.0/24` subnet must be routed into the `work` zone.
- The `work` zone should have all the necessary ports for **https** opened, but unencrypted **http** traffic should be filtered.

1. Verify that `firewalld` is enabled and running on your **serverX** system.

1.1.

```
[student@serverX ~]$ sudo systemctl status firewalld.service
```

Verify that the **Loaded** line ends in **enabled** and that the **Active** line specifies **running**. If this is not the case, enable and start the `firewalld.service` service using `systemctl`.

2. Install the `httpd` and `mod_ssl` packages.

2.1.

```
[student@serverX ~]$ sudo yum install httpd mod_ssl
```

3. Enable and start the `httpd.service` service.

- 3.1.

```
[student@serverX ~]$ sudo systemctl enable httpd.service
[student@serverX ~]$ sudo systemctl start httpd.service
```
- 4. Create the placeholder `/var/www/html/index.html` file with the contents `COFFEE!`.
 - 4.1.

```
[student@serverX ~]$ sudo vim /var/www/html/index.html
[student@serverX ~]$ cat /var/www/html/index.html
COFFEE!
```
- 5. Configure the `firewalld` daemon on `serverX` to route all traffic through the `dmz` zone by default.
 - 5.1.

```
[student@serverX ~]$ sudo firewall-cmd --set-default-zone=dmz
```
- 6. Configure the `firewalld` daemon on `serverX` to route all traffic from `172.25.X.0/24` through the `work` zone.
 - 6.1.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --zone=work --add-source=172.25.X.0/24
```
- 7. Open up incoming `https` traffic for the `work` zone.
 - 7.1.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --zone=work --add-service=https
```
- 8. Activate the changes to your firewall.
 - 8.1.

```
[student@serverX ~]$ sudo firewall-cmd --reload
```
- 9. Inspect the configuration of the running firewall on your `serverX` system.
 - 9.1.

```
[student@serverX ~]$ sudo firewall-cmd --get-default-zone
dmz
```
 - 9.2.

```
[student@serverX ~]$ sudo firewall-cmd --get-active-zones
dmz
  interfaces: eth0
work
  sources: 172.25.X.0/24
ROL
  sources: 172.25.0.252/32
```

□ 9.3.

```
[student@serverX ~]$ sudo firewall-cmd --zone=work --list-all
work
interfaces:
sources: 172.25.X.0/24
services: dhcpv6-client https ipp-client ssh
ports:
masquerade: no
forward-ports:
icmp-blocks:
rich rules:
```

- 10. From your **desktopX** system, use **curl** to test access to both `http://serverX.example.com` and `https://serverX.example.com`. The unencrypted connection should fail with a **No route to host** error message, while the encrypted connection should show your content.

Since **curl** does not trust the placeholder self-signed certificate on **serverX**, you will have to use the **-k** option to skip certificate validation.

□ 10.1.

```
[student@desktopX ~]$ curl http://serverX.example.com
curl: (7) Failed connect to serverX.example.com:80; No route to host
```

□ 10.2.

```
[student@desktopX ~]$ curl -k https://serverX.example.com
COFFEE!
```

Managing Rich Rules

Objectives

After completing this section, students should be able to configure more complex firewall configurations using firewalld's support for "rich language rules."

Rich rules concepts

Apart from the regular zones and services syntax that **firewalld** offers, administrators have two other options for adding firewall rules: *direct rules* and *rich rules*.

Direct rules

Direct rules allow an administrator to insert hand-coded **{ip, ip6, eb}tables** rules into the zones managed by **firewalld**. While powerful, and exposing features of the kernel **netfilter** subsystem not exposed through other means, these rules can be hard to manage. Direct rules also offer less flexibility than standard rules and rich rules. Configuring direct rules is not covered in this course, but documentation is available in the **firewall-cmd(1)** and **firewalld.direct(5)** man pages for those administrators who are already familiar with **{ip, ip6, eb}tables** syntax.

Unless explicitly inserted into a zone managed by **firewalld**, direct rules will be parsed before any **firewalld** rules are.

A short example of adding some direct rules to blacklist an IP range:

```
[root@serverX ~]# firewall-cmd --direct --permanent --add-chain ipv4 raw blacklist
[root@serverX ~]# firewall-cmd --direct --permanent --add-rule ipv4 raw PREROUTING 0 -s
192.168.0.0/24 -j blacklist
[root@serverX ~]# firewall-cmd --direct --permanent --add-rule ipv4 raw blacklist 0 -m
limit --limit 1/min -j LOG --log-prefix "blacklisted "
[root@serverX ~]# firewall-cmd --direct --permanent --add-rule ipv4 raw blacklist 1 -j
DROP
```

Rich rules

firewalld *rich rules* give administrators an expressive language in which to express custom firewall rules that are not covered by the basic **firewalld** syntax; for example, to only allow connections to a service from a single IP address, instead of all IP addresses routed through a zone.

Rich rules can be used to express basic allow/deny rules, but can also be used to configure logging, both to **syslog** and **auditd**, as well as port forwards, masquerading, and rate limiting.

The basic syntax of a rich rule can be expressed by the following block:

```
rule
  [source]
  [destination]
  service|port|protocol|icmp-block|masquerade|forward-port
  [log]
  [audit]
  [accept|reject|drop]
```


Almost every single element of a rule can take additional arguments in the form of **option=value**.



Note

For the full available syntax for rich rules, consult the **firewalld.richlanguage(5)** man page.

Rule ordering

Once multiple rules have been added to a zone (or the firewall in general), the ordering of rules can have a big effect on how the firewall behaves.

The basic ordering of rules inside a zone is the same for all zones:

1. Any port forwarding and masquerading rules set for that zone.
2. Any logging rules set for that zone.
3. Any allow rules set for that zone.
4. Any deny rules set for that zone.

In all cases, the first match will win. If a packet has not been matched by any rule in a zone, it will typically be denied, but zones might have a different default; for example, the **trusted** zone will **accept** any unmatched packet. Also, after matching a logging rule, a packet will continue to be processed as normal.

Direct rules are an exception. Most direct rules will be parsed before any other processing is done by **firewalld**, but the direct rule syntax allows an administrator to insert any rule they want anywhere in any zone.

Testing and debugging

To make testing and debugging easier, almost all rules can be added to the *runtime* configuration with a timeout. The moment the rule with a timeout is added to the firewall, the timer starts counting down for that rule. Once the timer for a rule has reached zero seconds, that rule is removed from the *runtime* configuration.

Using timeouts can be an incredibly useful tool while working on a remote firewalls, especially when testing more complicated rule sets. If a rule works, the administrator can add it again, but with the **--permanent** option (or at least without a timeout). If the rule does not work as intended, maybe even locking the administrator out of the system, it will be removed automatically, allowing the administrator to continue his or her work.

A timeout is added to a runtime rule by adding the option **--timeout=<TIMEINSECONDS>** to the end of the **firewall-cmd** that enables the rule.

Working with rich rules

firewall-cmd has four options for working with rich rules. All of these options can be used in combination with the regular **--permanent** or **--zone=<ZONE>** options.

Option	Explanation
--add-rich-rule='<RULE>'	Add <RULE> to the specified zone, or the default zone if no zone is specified.

Option	Explanation
<code>--remove-rich-rule=<RULE></code>	Remove <RULE> to the specified zone, or the default zone if no zone is specified.
<code>--query-rich-rule=<RULE></code>	Query if <RULE> has been added to the specified zone, or the default zone if no zone is specified. Returns 0 if the rule is present, otherwise 1 .
<code>--list-rich-rules</code>	Outputs all rich rules for the specified zone, or the default zone if no zone is specified.

Any configured rich rules are also shown in the output from `firewall-cmd --list-all` and `firewall-cmd --list-all-zones`.

Rich rules examples

Some examples of rich rules:

```
[root@serverX ~]# firewall-cmd --permanent --zone=classroom --add-rich-rule='rule
family=ipv4 source address=192.168.0.11/32 reject'
```

Reject all traffic from the IP address **192.168.0.11** in the **classroom** zone.

When using **source** or **destination** with an **address** option, the **family=** option of **rule** must be set to either **ipv4** or **ipv6**.

```
[root@serverX ~]# firewall-cmd --add-rich-rule='rule service name=ftp limit value=2/m
accept'
```

Allow two new connections to **ftp** per minute in the default zone.

Note that this change is only made in the *runtime* configuration.

```
[root@serverX ~]# firewall-cmd --permanent --add-rich-rule='rule protocol value=esp
drop'
```

Drop all incoming IPsec **esp** protocol packets from anywhere in the default zone.



Note

The difference between **reject** and **drop** lies in the fact that a **reject** will send back an ICMP packet detailing that, and why, a connection was rejected. A **drop** just drops the packet and does nothing else. Normally an administrator will want to use **reject** for friendly and neutral networks, and **drop** only for hostile networks.

```
[root@serverX ~]# firewall-cmd --permanent --zone=vnc --add-rich-rule='rule
family=ipv4 source address=192.168.1.0/24 port port=7900-7905 protocol=tcp accept'
```

Accept all TCP packets on ports **7900**, up to and including port **7905**, in the **vnc** zone for the **192.168.1.0/24** subnet.

Logging with rich rules

When debugging, or monitoring, a firewall, it can be useful to have a log of accepted or rejected connections. **firewalld** can accomplish this in two ways: by logging to **syslog**, or by sending messages to the kernel **audit** subsystem, managed by **auditd**.

In both cases, logging can be *rate limited*. Rate limiting ensures that system log files do not fill up with messages at a rate such that the system cannot keep up, or fills all its disk space.

The basic syntax for logging to **syslog** using rich rules is:

```
log [prefix="<PREFIX TEXT>" [level=<LOGLEVEL>] [limit value="<RATE/DURATION>"]
```

Where **<LOGLEVEL>** is one of **emerg**, **alert**, **crit**, **error**, **warning**, **notice**, **info**, or **debug**.

<DURATION> can be one of **s** for seconds, **m** for minutes, **h** for hours, or **d** for days. For example, **limit value=3/m** will limit the log messages to a maximum of three per minute.

The basic syntax for logging to the audit subsystem is:

```
audit [limit value="<RATE/DURATION>"]
```

Rate limiting is configured in the same way as for **syslog** logging.

Logging examples

Some examples of logging using rich rules:

```
[root@serverX ~]# firewall-cmd --permanent --zone=work --add-rich-rule='rule service
name="ssh" log prefix="ssh " level="notice" limit value="3/m" accept
```

Accept new connections to **ssh** from the **work** zone, log new connections to **syslog** at the **notice** level, and with a maximum of three message per minute.

```
[root@serverX ~]# firewall-cmd --add-rich-rule='rule family=ipv6 source
address="2001:db8::/64" service name="dns" audit limit value="1/h" reject' --
timeout=300
```

New IPv6 connections from the subnet **2001:db8::/64** in the default zone to DNS are rejected for the next five minutes, and rejected connections are logged to the **audit** system with a maximum of one message per hour.

R

References

firewalld.richlanguage(5), **firewall-cmd(1)**, and **firewalld.direct(5)** man pages

Practice: Writing Custom Rules

Guided exercise

In this lab, you will configure your **serverX** system to allow connections to a (new) **http** service, but only from **desktopX**, and with a rate-limited log message.

Resources:

Machines:	desktopX and serverX
------------------	------------------------------------

Outcomes:

Custom firewall rules that configure rate-limited logging for specific connections.

Before you begin...

- Reset your **desktopX** system.
- Reset your **serverX** system.

Your company is running a trial that includes starting a web server on **serverX**, but for the duration of the trial, only **desktopX** should be able to connect. Since this could potentially generate many log entries, this logging should be limited to a maximum of three messages per second, and all log messages should be prefixed with the message "**NEW HTTP**".

It has been decided that you, the IT Rock Star, will implement this using **firewalld rich rules**.

1. First install, start, and enable **httpd**.

- 1.1.

```
[student@serverX ~]$ sudo yum install httpd
```

- 1.2.

```
[student@serverX ~]$ sudo systemctl start httpd.service
```

- 1.3.

```
[student@serverX ~]$ sudo systemctl enable httpd.service
```

2. Configure a firewall rule in the default zone that allows traffic to **http** only from your **desktopX** system. This traffic should be logged, but with a maximum of three new connections per second.

- 2.1. Permanently create the new firewall rule.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --add-rich-rule='rule
family=ipv4 source address=172.25.X.10/32 service name="http" log
level=notice prefix="NEW HTTP " limit value="3/s" accept'
```

- 2.2. Activate the changes to your firewall.

```
[student@serverX ~]$ sudo firewall-cmd --reload
```

- 2.3. On your **serverX** system, use **tail -f** to view the additions to **/var/log/messages** in real time.

```
[student@serverX ~]$ sudo tail -f /var/log/messages
```

- 2.4. From your **desktopX** system, use **curl** to connect to the **httpd** service running on **serverX**.

```
[student@desktopX ~]$ curl http://serverX.example.com
```

- 2.5. Inspect the output of your running **tail** command on **serverX**. You should see a message for the new connection like this:

```
May  9 08:04:11 serverX kernel: NEW HTTP IN=eth0 OUT= MAC=...  
SRC=172.25.X.10 DST=172.25.X.11 LEN=60....
```

Masquerading and Port Forwarding

Objectives

After completing this section, students should be able to describe and implement Network Address Translation (NAT).

Network Address Translation (NAT)

`firewalld` supports two types of *Network Address Translation* (NAT): *masquerading* and *port forwarding*. Both can be configured on a basic level with regular `firewall-cmd` rules, and more advanced forwarding configurations can be accomplished with rich rules. Both forms of NAT modify certain aspects of a packet, like the source or destination, before sending it on.

Masquerading

With *masquerading*, a system will forward packets that are not directly addressed to itself to the intended recipient, while changing the *source address* of the packets that go through to its own public IP address. When answers to those packets come in, the firewall will then modify the destination address to the address of the original host, and send the packet on. This is usually used on the edge of a network to provide Internet access to an internal network. Masquerading is a form of *Network Address Translation* (NAT).



Important

Masquerading can only be used with **IPv4**, not with **IPv6**.

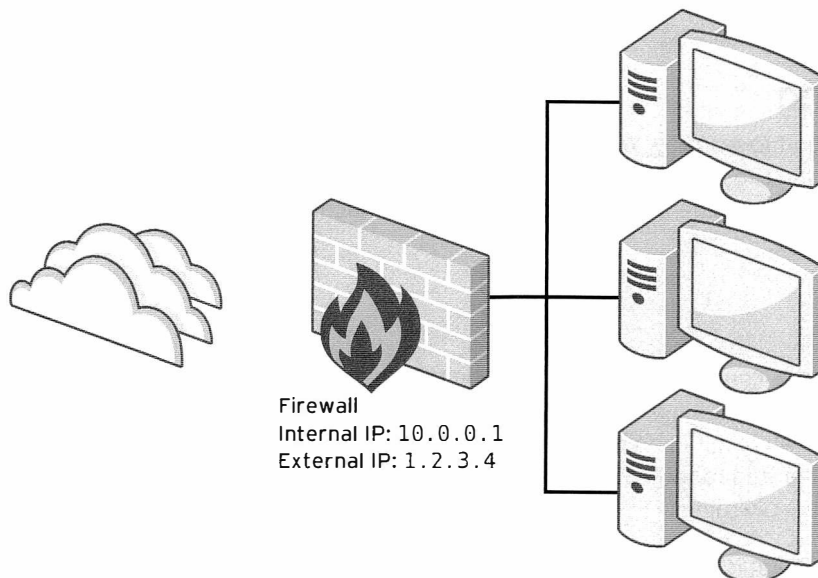


Figure 4.1: A sample network layout for NAT

An example of how masquerading works based on the network layout described in *Figure 4.1: A sample network layout for NAT*.

1. One of the machines behind the firewall sends a packet to an address outside of the local network. The packet has a source address of **10.0.0.100** (the address of the machine), and a destination address of **2.17.39.214**.
2. Since the destination address is not on the local subnet, the packet will be routed to the default gateway configured on the source machine; in this case, **10.0.0.1**, the IP address of the firewall.
3. The firewall accepts the packet, changes the source address to **1.2.3.4** (the external IP for the firewall), stores a reference to this connection in its connection state table, then passes it to a router on the Internet based on its routing table.
4. An answer to the packet comes back from the Internet. The router looks up the connection in its connection state table, then changes the destination address to **10.0.0.100** (the original sender), and passes the packet on.
5. The original sender receives the answer to its request.

Configuring masquerading

To configure masquerading for a zone with regular **firewall-cmd** commands, use the following syntax:

```
[root@serverX ~]# firewall-cmd --permanent --zone=<ZONE> --add-masquerade
```

This will masquerade any packets sent to the firewall from clients defined in the sources for that zone (both interfaces and subnets) that are not addressed to the firewall itself.

To gain more control over what clients will be masqueraded, a rich rule can be used as well.

```
[root@serverX ~]# firewall-cmd --permanent --zone=<ZONE> --add-rich-rule='rule
family=ipv4 source address=192.168.0.0/24 masquerade'
```

Port forwarding

Another form of NAT is *port forwarding*. With port forwarding, traffic to a single port is forwarded either to a different port on the same machine, or to a port on a different machine. This mechanism is typically used to “hide” a server behind another machine, or to provide access to a service on an alternate port.



Important

When a port forward is configured to forward packets to a different machine, any replies from that machine will normally be sent directly to the original client from that machine. Since this will result in an invalid connection on most configurations, the machine that is forwarded to will have to be masqueraded through the firewall that performed the port forwarding.

A common configuration is to forward a port from the firewall machine to a machine that is already masqueraded behind the firewall.

An example of a port forward based on the network layout described in *Figure 4.1: A sample network layout for NAT*. Assume that the machine with the IP address **10.0.0.100** behind the

firewall is running a web server on port **8080/TCP**, and that the firewall is configured to forward traffic coming in on port **80/TCP** on its external interface to port **8080/TCP** on that machine.

1. A client from the Internet sends a packet to port **80/TCP** on the external interface of the firewall.
2. The firewall changes the destination address and port of this packet to **10.0.0.100** and **8080/TCP** and forwards it on. The source address and port remain unchanged.
3. The machine behind the firewall sends a response to this packet. Since this machine is being masqueraded (and the firewall is configured as the default gateway), this packet is sent to the original client, appearing to come from the external interface on the firewall.

Configuring port forwarding

To configure port forwarding with regular **firewall-cmd** commands, use the following syntax:

```
[root@serverX ~]# firewall-cmd --permanent --zone=<ZONE> --add-forward-
port=port=<PORTNUMBER>;proto=<PROTOCOL>[:toport=<PORTNUMBER>][:toaddr=<IPADDR>]
```

Both the **toport=** and **toaddr=** parts are optional, but *at least* one of those two will need to be specified.

As an example, the following command will forward incoming connections on port **513/TCP** on the firewall to port **132/TCP** on the machine with the IP address **192.168.0.254** for clients from the **public** zone:

```
[root@serverX ~]# firewall-cmd --permanent --zone=public --add-forward-
port=port=513;proto=tcp;toport=132;toaddr=192.168.0.254
```

To gain more control over port forwarding rules, the following syntax can be used with rich rules:

```
forward-port port=<PORTNUM> protocol=tcp|udp [to-port=<PORTNUM>] [to-addr=<ADDRESS>]
```

An example that uses rich rules to forward traffic from **192.168.0.0/26** in the **work** zone to port **80/TCP** to port **8080/TCP** on the firewall machine itself:

```
[root@serverX ~]# firewall-cmd --permanent --zone=work --add-rich-rule='rule family=ipv4
source address=192.168.0.0/26 forward-port port=80 protocol=tcp to-port=8080'
```

R

References

firewalld.richlanguage(5) and **firewall-cmd**(1) man pages

Practice: Forwarding a Port

Guided exercise

In this lab, you will configure your **serverX** system to forward a request to port **443/tcp** from **desktopX** to port **22/tcp**.

Resources:	
Machines:	desktopX and serverX

Outcomes:

Custom firewall rules that configure a port forward.

Before you begin...

- Reset your **desktopX** system.
- Reset your **serverX** system.

Your company is running a trial for a new bastion host. As part of this trial, your **desktopX** should be able to connect to the SSH daemon on your **serverX** system on port **443/tcp**. Since this is purely a trial, you do not wish to bind **sshd** to that port directly, and only your **desktopX** should be able to connect using port **443/tcp**.

It has been decided that you, the chosen one, will implement this using **firewalld rich rules**.

- 1. Configure the firewall on **serverX** to forward port **443/tcp** to **22/tcp**, but only for your **desktopX** machine. The IP address of your **desktopX** machine is **172.25.X.10**.
 - 1.1. Permanently add the port forwarding firewall rule on **serverX**.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --add-rich-rule
'rule family=ipv4 source address=172.25.X.10/32 forward-port port=443
protocol=tcp to-port=22'
```

- 1.2. Reload the firewall configuration to activate your changes.

```
[student@serverX ~]$ sudo firewall-cmd --reload
```

- 1.3. Test if **sshd** is now available on port **443/tcp** from your **desktopX** system.

```
[student@desktopX ~]$ ssh -p 443 serverX.example.com
The authenticity of host '[serverX.example.com]:443 ([172.25.X.11]:443)'
can't be established.
ECDSA key fingerprint is XX:XX:XX:XX:XX:XX:XX:XX:XX:XX:XX:XX:XX:XX:XX.
Are you sure you want to continue connecting (yes/no)? yes
student@serverX.example.com's password: student
```

Managing SELinux Port Labeling

Objectives

After completing this section, students should be able to ensure network ports have the correct SELinux type so that services are able to bind to them.

SELinux port labeling

SELinux does more than just file and process labeling. Network traffic is also tightly enforced by the SELinux policy. One of the methods that SELinux uses for controlling network traffic is labeling network ports; for example, in the **targeted** policy, port **22/TCP** has the label **ssh_port_t** associated with it.

Whenever a process wants to listen on a port, SELinux will check to see if the label associated with that process (the domain) is allowed to bind that port label. This can stop a rogue service from taking over ports otherwise used by other (legitimate) network services.

Managing SELinux port labeling

Whenever an administrator decides to run a service on a nonstandard port, there is a high chance that SELinux port labels will need to be updated. In some cases, the **targeted** policy has already labeled the port with a type that can be used; for example, since port **8008/TCP** is often used for web applications, that port is already labeled with **http_port_t**, the default port type for the web server.

Listing port labels

To get an overview of all the current port label assignments, administrators can use the **port** subcommand of the **semanage** command. The **-l** option will list all current assignments, in the form:

```
port_label_t tcp|udp comma, separated, list, of, ports
```

To only view local changes to the default policy, administrators can add the **-C** option to this command.

Example output:

```
[root@serverX ~]# semanage port -l
...
http_cache_port_t      tcp    8080, 8118, 8123, 10001-10010
http_cache_port_t      udp    3130
http_port_t            tcp    80, 81, 443, 488, 8008, 8009, 8443, 9000
...
```

Note that a port label can appear twice in the output, once for **TCP** and once for **UDP**.



Note

A graphical tool, **system-config-selinux**, is also available for administrators who prefer to work with GUI tools. This tool is part of the *policycoreutils-gui* package.

Managing port labels

semanage can also be used to assign new port labels, remove port labels, or modify existing ones.



Important

Only local modifications can be removed or modified. To allow a service to bind to a port label normally associated with another service, a small policy module must be written. Writing and generating policy modules falls outside the scope of this course. It is also not possible to remove a port label from the policy without overriding the policy module that provided that label.

To add a port to an existing port label (type), use the following syntax:

```
[root@serverX ~]# semanage port -a -t port_label -p tcp|udp PORTNUMBER
```

For example, to allow a **gopher** service to listen on port **71/TCP**:

```
[root@serverX ~]# semanage port -a -t gopher_port_t -p tcp 71
```



Note

The **targeted** policy ships with a large number of port types. Per-service documentation on SELinux types, Booleans, and port types can be found in the service-specific SELinux man pages found in the *selinux-policy-devel* package.

If these man pages are not yet installed on your system, follow this procedure:

```
[root@serverX ~]# yum -y install selinux-policy-devel  
[root@serverX ~]# mandb  
[root@serverX ~]# man -k _selinux
```

Removing port labels

The syntax for removing a custom port label is the same as the syntax for adding a port label, but instead of using the **-a** option (for Add), the **-d** option (for Delete) is used.

For example, to remove the binding of port **71/TCP** to **gopher_port_t**:

```
[root@serverX ~]# semanage port -d -t gopher_port_t -p tcp 71
```

Modifying port bindings

If an administrator has accidentally assigned the wrong type to a port, or requirements have changed, it's possible to modify the label associated with a port. This is a more efficient process than removing the old binding and adding a new one. Modifications require the **-m** option.

For example, to modify port **71/TCP** from **gopher_port_t** to **http_port_t**, an administrator can use the following command:

```
[root@serverX ~]# semanage port -m -t http_port_t -p tcp 71
```



References

semanage(8), **semanage-port(8)**, ***_selinux(8)**, and **system-config-selinux(8)** man pages

Practice: Managing SELinux Port Labeling

Guided exercise

In this lab, you will configure your **serverX** system to allow **http** access on a nonstandard port.

Resources:

Machines:	desktopX and serverX
------------------	------------------------------------

Outcomes:

A web server running on **serverX** successfully serving content on a nonstandard port.

Before you begin...

- Reset your **desktopX** system.
- Reset your **serverX** system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab selinuxport setup
```

Your organization is deploying a new custom web application. Unfortunately for you, the web application is running on a nonstandard port; in this case, **82/TCP**.

One of your developers has already configured the application on your **serverX**. But, not being a rock star system administrator, he failed in getting the web server to start successfully. Your mission, if you choose to accept it, is to get the **httpd.service** service on **serverX** successfully started, and serving out content to your **desktopX** system over port **82/TCP**.

1. Start by restarting the **httpd.service**.

- 1.1.

```
[student@serverX ~]$ sudo systemctl restart httpd.service
Job for httpd.service failed. See 'systemctl status httpd.service' and
'journalctl -xn' for details
```

- 1.2. View the output from **systemctl status -l httpd.service**.

```
[student@serverX ~]$ sudo systemctl status -l httpd.service
...
Permission denied: AH00072: make_sock: could not bind to address
0.0.0.0:82
...
```

- 1.3. Check if SELinux is blocking **httpd** from binding to port **82/TCP**.

```
[student@serverX ~]$ sudo sealert -a /var/log/audit/audit.log
```

2. Configure SELinux to allow **httpd** to bind to port **82/TCP**, then restart the **httpd.service** service.

- 2.1. Use **semanage** to find an appropriate port type for port **82/TCP**.

```
[student@serverX ~]$ sudo semanage port -l | grep http
```

`http_port_t` seems promising, since it is what the normal `http` port (`80/TCP`) is also assigned to.

- 2.2. Assign port `82/TCP` the `http_port_t` type.

```
[student@serverX ~]$ sudo semanage port -a -t http_port_t -p tcp 82
```

- 2.3. Restart the `httpd.service` service.

```
[student@serverX ~]$ sudo systemctl restart httpd.service
```

- 3. Check if you can now access the web server running on port `82/TCP`.

- 3.1.

```
[student@serverX ~]$ curl http://serverX.example.com:82  
Hello
```

- 4. Check if you can access the new web service from your `desktopX` system.

- 4.1.

```
[student@desktopX ~]$ curl http://serverX.example.com:82  
curl: (7) Failed to connect to serverX.example.com:82; No route to host
```

That error means you still can't connect from `desktopX`. Take a minute to think up some probable causes for this failure.

- 5. On your `serverX` system, open up port `82/TCP` on your firewall.
 - 5.1. Open port `82/TCP` in the permanent configuration for the default zone on the firewall on `serverX`.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --add-port=82/tcp
```

- 5.2. Activate your firewall changes on `serverX`.

```
[student@serverX ~]$ sudo firewall-cmd --reload
```

- 6. Check if you can now access the new web service from your `desktopX` system.

- 6.1.

```
[student@desktopX ~]$ curl http://serverX.example.com:82  
Hello
```

- 7. On your `desktopX` system, run the command `lab selinuxport grade` to verify your work.

□ 7.1.

```
[student@desktopX ~]$ lab selinuxport grade
Testing http access over port 82... PASS
```

Lab: Network Port Security

Performance checklist

In this lab, you will configure a custom firewall zone, service, and a custom SELinux port type.

Resources:	
Machines:	desktopX and serverX

Outcomes:

An SSHD service listening on both the regular port **22/TCP**, as well as the new nonstandard port **999/TCP**, protected through a custom firewall zone and a custom firewall service.

Before you begin...

- Reset your **desktopX** system.
- Reset your **serverX** system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab custom-sshd setup
```

Since some of your co-workers sometimes have to use your **serverX** machine from a network that blocks outgoing access to port **22/TCP**, you have decided to bind your **sshd** daemon to the nonstandard port **999/TCP**. The **sshd** daemon should also still be available over the standard port **22/TCP**. The **lab custom-sshd setup** script has already configured the **sshd** daemon to listen on these ports, but has not yet performed any additional configuration that might be necessary.

After ensuring **sshd** listens on these ports, you will need to update the firewall on **serverX** as well. Use the following criteria when configuring the firewall:

- All connections from the **172.25.X.0/24** network should be routed through the **work** zone.
- In addition to the regular **ssh** port, port **999/TCP** should also be available in the **work** zone.

After you have completed your work, you can run the command **lab custom-sshd grade** on your **desktopX** machine to verify your work.

1. Verify that **sshd** is indeed listening on both port **22/TCP** and port **999/TCP**. If not, investigate and fix.
2. Route all traffic from the **172.25.X.0/24** subnet through the **work** zone.
3. Open port **999/TCP** in the **work** zone.
4. From **desktopX**, verify your configuration by running **lab custom-sshd grade**.

Solution

In this lab, you will configure a custom firewall zone, service, and a custom SELinux port type.

Resources:	
Machines:	desktopX and serverX

Outcomes:

An SSHD service listening on both the regular port **22/TCP**, as well as the new nonstandard port **999/TCP**, protected through a custom firewall zone and a custom firewall service.

Before you begin...

- Reset your **desktopX** system.
- Reset your **serverX** system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab custom-sshd setup
```

Since some of your co-workers sometimes have to use your **serverX** machine from a network that blocks outgoing access to port **22/TCP**, you have decided to bind your **sshd** daemon to the nonstandard port **999/TCP**. The **sshd** daemon should also still be available over the standard port **22/TCP**. The **lab custom-sshd setup** script has already configured the **sshd** daemon to listen on these ports, but has not yet performed any additional configuration that might be necessary.

After ensuring **sshd** listens on these ports, you will need to update the firewall on **serverX** as well. Use the following criteria when configuring the firewall:

- All connections from the **172.25.X.0/24** network should be routed through the **work** zone.
- In addition to the regular **ssh** port, port **999/TCP** should also be available in the **work** zone.

After you have completed your work, you can run the command **lab custom-sshd grade** on your **desktopX** machine to verify your work.

1. Verify that **sshd** is indeed listening on both port **22/TCP** and port **999/TCP**. If not, investigate and fix.
 - 1.1. Use **ss** to find all the ports where **sshd** is listening.

```
[student@serverX ~]$ sudo ss -tulpn | grep sshd
tcp    LISTEN  0      128          *:22          *:*
      users:(("sshd", 4864, 3))
tcp    LISTEN  0      128          :::22         :::*
      users:(("sshd", 4864, 4))
```

- 1.2. The previous output indicates that **sshd** is only bound to port **22/TCP**. Restart the **sshd.service** service and check the status output.

```
[student@serverX ~]$ sudo systemctl restart sshd.service
[student@serverX ~]$ sudo systemctl status -l sshd.service
...
May 12 12:41:37 serverX.example.com sshd[23036]: error: Bind to port 999 on
0.0.0.0 failed: Permission denied.
...
```

Depending on the speed and configuration of your system the previous command might also include some **sealert** output. If an SELinux was caused by a service this information typically shows up with some delay in the output of **systemctl status**, but a lack of **sealert** output is not a guarantee that this is not an SELinux issue.

- 1.3. That **Permission denied** while binding to port **999/TCP** is most likely an SELinux denial. Use **sealert** to find any SELinux messages related to this.

```
[student@serverX ~]$ sudo sealert -a /var/log/audit/audit.log
...
SELinux is preventing /usr/sbin/sshd from name_bind access on the tcp socket.
...
```

- 1.4. Use the advice from **sealert** to allow **sshd** to bind to port **999/TCP**.

```
[student@serverX ~]$ sudo semanage port -a -t ssh_port_t -p tcp 999
```

- 1.5. Restart **sshd.service** and verify if it is now bound to port **999/TCP** as well.

```
[student@serverX ~]$ sudo systemctl restart sshd.service
[student@serverX ~]$ sudo ss -tulpn | grep sshd
...
tcp    LISTEN  0      128                *:999              *.*
users:(("sshd",4864,3))
...
```

2. Route all traffic from the **172.25.X.0/24** subnet through the **work** zone.
 - 2.1. Permanently route all traffic from the **172.25.X.0/24** subnet through the **work** zone.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --zone=work --add-source=172.25.X.0/24
```

3. Open port **999/TCP** in the **work** zone.

- 3.1. Add the port.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --zone=work --add-port=999/tcp
```

- 3.2. Reload the firewall on **serverX** to activate your changes.

```
[student@serverX ~]$ sudo firewall-cmd --reload
```

4. From **desktopX**, verify your configuration by running **lab custom-sshd grade**.

4.1.

```
[student@desktopX ~]$ lab custom-sshd grade
...
Overall result: PASS
```

Summary

Managing Firewalld

In this section, students reviewed firewalld concepts and management commands covered in previous courses.

Managing Rich Rules

In this section, students learned how to configure more complex firewall configurations using firewalld's support for "rich language rules."

Masquerading and Port Forwarding

In this section, students learned how to describe and implement Network Address Translation (NAT).

Managing SELinux Port Labeling

In this section, students learned how to ensure network ports have the correct SELinux type so that services are able to bind to them.



CHAPTER 5

MANAGING DNS FOR SERVERS

Overview	
Goal	To set and verify correct DNS records for systems and configure secure caching DNS name service.
Objectives	<ul style="list-style-type: none">• Explain how DNS is used to resolve names and addresses and the purpose of key DNS resource records.• Configure <i>unbound</i> to act as a secure local caching nameserver.• Describe common DNS configuration problems and how to identify and resolve them.
Sections	<ul style="list-style-type: none">• DNS Concepts (and Practice)• Configuring a Caching Nameserver (and Practice)• DNS Troubleshooting (and Practice)
Lab	<ul style="list-style-type: none">• Managing DNS for Servers

DNS Concepts

Objectives

After completing this section, students should be able to:

- Explain the hierarchical structure of the Domain Name System (DNS).
- Differentiate between domains, subdomains, and zones.
- Identify the differences between different resource record types.

The Domain Name System

The *Domain Name System (DNS)* is a hierarchical naming system that serves as a directory of networked hosts and resources. Information in the directory maps network names to data and is maintained in logical entries known as resource records. The DNS hierarchy begins with the *root* domain "." at the top and branches downward to multiple next-level domains.

Each level of the DNS hierarchy is delineated by the "." in domain names, with "." as the top level. Domains such as **com**, **net**, and **org** occupy the second level of the hierarchy and domains such as **example.com** and **redhat.com** occupy the third level and so on.

When working with DNS, it is important to clarify some of the common terms used to refer to the structure of the DNS hierarchy, such as **domain**, **subdomain**, and **zone**.

Domain

A **domain** is a collection of resource records that ends in a common name and represents an entire subtree of the DNS name space, such as **example.com**. The largest possible domain is the *root* domain, ".", which includes the whole DNS namespace.

A *top-level domain (TLD)* is a domain that has only one component. *Generic TLDs (gTLDs)* were originally organized by theme, and include **.com**, **.edu**, **.net**, etc. *Country code TLDs (ccTLDs)* are organized on a national basis, and include **.us**, **.uk**, **.cn**, **.ru**, etc.

Subdomain

A **subdomain** is a domain that is a subtree of another domain. This term is used when discussing the relationship of two domains to each other. For example, **lab.example.com** is a subdomain of **example.com**.

Zone

A **zone** is the portion of a domain for which a particular nameserver is directly responsible, or authoritative. This may be an entire domain, or just part of a domain with some or all of its subdomains delegated to other nameserver(s).

Anatomy of DNS lookups

When a system needs to perform name resolution using a DNS server, it begins by sending queries to the servers listed in **/etc/resolv.conf** in order, until it gets a response or runs out of servers. The **host** or **dig** commands can be used to manually look up DNS names.

Local authoritative data

When the query arrives at a DNS server, the server first determines whether the information being queried resides in a zone that it is authoritative for. If the server is an authority for the

zone that the name or address being queried belongs to, then the server responds to the client with the information contained in its local zone file. This type of response is referred to as an *authoritative answer* (*aa*), since the server providing the response is authoritative for the data provided. Authoritative answers from a nameserver have the **aa** flag turned on in the header of the DNS response.

Local cached non-authoritative data

If the DNS server is not an authority for the record in question, but has recently obtained the record to answer a previous query, it may still have a copy of the record in its cache. The cache is where answers to queries are stored for a specified time, determined by a value contained in every resource record response called the *Time To Live* (*TTL*). If an answer exists in the server's cache, it is provided to the client. This answer will not have the **aa** flag set, since the server is not authoritative for the data being provided.

Remote non-authoritative data via recursion

If the DNS server is not authoritative for the name being queried, and it does not possess the record in its cache, it will then attempt to retrieve the record via an iterative process known as recursion. A DNS server with an empty cache begins the recursion process by querying one of the root nameservers by IP address retrieved from its local, pre-populated *root hints file*. The root nameserver will then likely respond with a referral, which indicates the nameservers that are authoritative for the TLD that contains the name being queried.

Upon receiving the referral, the DNS server will then perform another iterative query to the TLD authoritative nameserver it was referred to. Depending on whether there are further remaining delegations in the name being queried, this authoritative nameserver will either send an authoritative answer or yet another referral. This continues until an authoritative server is reached and responds with an authoritative answer.

The final answer, along with all the intermediate answers obtained prior to it, are cached by the DNS server to improve performance. If during a lookup for **www.example.com** the DNS server finds out that the **example.com** zone has authoritative nameservers, it will query those servers directly for any future queries for information in the **example.com** zone, rather than starting recursion again at the root nameservers.

DNS resource records

DNS *resource records* (*RRs*) are entries in a DNS zone that specify information about a particular name or object in the zone. A resource record contains a **type**, a **TTL**, a **class**, and **data** elements organized in the following format:

<i>owner-name</i>	<i>TTL</i>	<i>class</i>	<i>type</i>	<i>data</i>
www.example.com.	300	IN	A	192.168.1.10

Resource Record Fields

Field name	Content
owner - name	The name for this resource record.
TTL	The <i>Time To Live</i> of the resource record in seconds. This specifies how long this resource record should be cached by DNS resolvers.
class	The "class" of the record, almost always IN ("Internet").

Field name	Content
type	The type indicates the sort of information stored by this record. For example, an A record maps a host name to an IPv4 address.
data	The data stored by this record. The exact format varies by record type.

There are a number of important resource record types:

A (IPv4 address) records

An **A** resource record maps a host name to an IPv4 address.

```
[student@serverX ~]$ host -v -t A example.com
Trying "example.com"
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 22681
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 2

;; QUESTION SECTION:
;example.com. IN A

;; ANSWER SECTION:
example.com. 86400 IN A 172.25.254.254

Received 96 bytes from 172.25.254.254#53 in 1 ms
```

AAAA (IPv6 address) records

An **AAAA** resource record ("quad-A" record) maps a host name to an IPv6 address.

```
[student@serverX ~]$ host -v -t AAAA a.root-servers.net
Trying "a.root-servers.net"
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 18194
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 13, ADDITIONAL: 12

;; QUESTION SECTION:
;a.root-servers.net. IN AAAA

;; ANSWER SECTION:
a.root-servers.net. 604800 IN AAAA 2001:503:ba3e::2:30

Received 64 bytes from 172.25.254.254#53 in 78 ms
```

CNAME (canonical name) records

A **CNAME** resource record aliases one name to another name (the *canonical name*), which should have **A** or **AAAA** records.

When a DNS resolver receives a **CNAME** record in response to a query, it will reissue the query using the canonical name instead of the original name.

The data field of **CNAME** records can point to a name anywhere in DNS, whether internal or external to the zone:

```
www-dev.example.com. IN CNAME lab.example.com.
www.example.com.    IN CNAME www.redhat.com.
```

CNAME records are useful, but should be used with some care. In general, pointing a **CNAME** records to other **CNAME** records should be avoided for efficiency and fragility reasons and to avoid creating a **CNAME** loop by accident. The chain of **CNAME** record must end in **A** and/or **AAAA**

records. Note that there are legitimate uses for **CNAME** chains when using Content Delivery Networks (CDNs) to improve the speed and reliability of data delivery over the Internet. Likewise, **NS** and **MX** records must not be pointed at **CNAME** records but at names with **A** and/or **AAAA** records.

```
[student@serverX ~]$ host -v -t A ipa-ca.server0.example.com
Trying "ipa-ca.server0.example.com"
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 11931
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 1, ADDITIONAL: 2

;; QUESTION SECTION:
;ipa-ca.server0.example.com. IN A

;; ANSWER SECTION:
ipa-ca.server0.example.com. 86400 IN CNAME server0.example.com.
server0.example.com. 86400 IN A 172.25.0.11

Received 125 bytes from 172.25.254.254#53 in 1 ms
```

PTR (pointer) records

A **PTR** record maps IPv4 or IPv6 addresses to a host name. They are used for *reverse DNS resolution*.

PTR records code the IP address in a special format that acts like a host name. For IPv4 addresses, the address is reversed, most specific part first, and the result is treated as a host in a subdomain of the special domain in-addr.arpa. For IPv6 addresses, the address is split into subdomains on nibble boundaries (every hexadecimal digit) and set up as a subdomain of the special domain ip6.arpa, as seen in the following example. While this syntax may seem strange, it makes it simpler for DNS administrators to delegate responsibility for ranges of addresses to other DNS administrators.

```
[student@serverX ~]$ host -v -t PTR 172.25.0.10
Trying "10.0.25.172.in-addr.arpa"
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 36389
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 2

;; QUESTION SECTION:
;10.0.25.172.in-addr.arpa. IN PTR

;; ANSWER SECTION:
10.0.25.172.in-addr.arpa. 86400 IN PTR desktop0.example.com.

Received 127 bytes from 172.25.254.254#53 in 2 ms
```

```
[student@serverX ~]$ host -v -t PTR 2001:503:ba3e::2:30
Trying "0.3.0.0.2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.e.3.a.b.3.0.5.0.1.0.0.2.ip6.arpa"
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 32138
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;0.3.0.0.2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.e.3.a.b.3.0.5.0.1.0.0.2.ip6.arpa. IN PTR

;; ANSWER SECTION:
0.3.0.0.2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.e.3.a.b.3.0.5.0.1.0.0.2.ip6.arpa. 86400 IN PTR
a.root-servers.net.

Received 122 bytes from 172.25.254.254#53 in 174 ms
```

NS (name server) records

An **NS** record maps a domain name to a DNS name server which is authoritative for its DNS zone.

Every public authoritative name server for the zone must have an **NS** record.

```
[student@serverX ~]$ host -v -t NS example.com
Trying "example.com"
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 29362
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 2

;; QUESTION SECTION:
;example.com. IN NS

;; ANSWER SECTION:
example.com. 86400 IN NS classroom.example.com.

Received 80 bytes from 172.25.254.254#53 in 0 ms
```

SOA (start of authority) records

An **SOA** record provides information about how a DNS zone works.

There will be exactly one **SOA** record for a zone. It specifies which of the zone's name servers is the primary one (the *master*), information on how secondary (*slave*) name servers should update their copy of the information, and the zone's management contact. Its data field contains the following elements:

SOA record data elements

Data element	Content
Master nameserver	The host name of the nameserver which is the original source of domain information, and which may accept dynamic DNS updates if the zone supports them.
RNAME	The email address of the person responsible for the DNS zone (the <i>hostmaster</i>). The @ in the email address is replaced with a "." in the RNAME. For example, an email address of hostmaster@example.com is written as hostmaster.example.com .
Serial number	The version number of the zone, which is increased when there is any change to zone records.
Refresh	How frequently the slave servers should check for zone updates, in seconds.
Retry	How long a slave server should wait before retrying a failed refresh attempt, in seconds.
Expiry	If refreshes have been failing, how long a slave server should wait before it stops using its old copy of the zone to respond to queries, in seconds.
Minimum	If a resolver looks up a name and it does not exist (gets a <i>nonexistent domain (NXDOMAIN)</i> response), how long it should cache the information that the record does not exist, in seconds.

```
[student@serverX ~]$ host -v -t SOA example.com
Trying "example.com"
;; -->HEADER<<- opcode: QUERY, status: NOERROR, id: 58434
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;example.com. IN SOA

;; ANSWER SECTION:
example.com. 86400 IN SOA classroom.example.com. root.classroom.example.com. 2013091600
3600 300 604800 60

Received 121 bytes from 172.25.254.254#53 in 0 ms
```

MX (mail exchange) records

An **MX** record maps a domain name to a *mail exchange* which will accept email for that name.

The data for this record type is a preference number (lowest preferred) used to determine the order in which to pick between multiple **MX** records, and a host name for a mail exchange for that name.

```
[student@serverX ~]$ host -v -t MX example.com
Trying "example.com"
;; -->HEADER<<- opcode: QUERY, status: NOERROR, id: 47187
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 2

;; QUESTION SECTION:
;example.com. IN MX

;; ANSWER SECTION:
example.com. 86400 IN MX 10 classroom.example.com.

Received 96 bytes from 172.25.254.254#53 in 0 ms
```

TXT (text) records

A **TXT** record is used to map a name to arbitrary human-readable text.

TXT records are commonly used to supply data used by *Sender Policy Framework (SPF)*, *DomainKeys Identified Mail (DKIM)*, *Domain-based Message Authentication, Reporting and Conformance (DMARC)*, and so on.

```
[student@serverX ~]$ host -v -t TXT lwn.net
Trying "lwn.net"
;; -->HEADER<<- opcode: QUERY, status: NOERROR, id: 41137
;; flags: qr rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;lwn.net. IN TXT

;; ANSWER SECTION:
lwn.net. 28619 IN TXT "v=spf1 ip4:72.51.34.34 ip4:70.33.254.29 -all"

Received 638 bytes from 192.168.2.11#53 in 74 ms
```

SRV (service) records

An **SRV** record is used to locate the hosts which support a particular service for a domain.

Using a domain name formatted to include a service and a protocol name, `_service._protocol.domainname`, **SRV** records provide the names of the hosts that provide that service for the domain, as well as the port number that the service listens on. **SRV** records also include **priority** and **weight** values to indicate the order in which hosts should be used when multiple hosts are available for a particular service.

This example **SRV** record indicates that the `server0.example.com` domain provides the **LDAP** service using **TCP** on port **389** on host `server0.example.com` with a priority of **0** and a weighting of **100**.

```
[student@serverX ~]$ host -v -t SRV _ldap._tcp.server0.example.com
Trying "_ldap._tcp.server0.example.com"
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 35665
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 3

;; QUESTION SECTION:
;_ldap._tcp.server0.example.com. IN SRV

;; ANSWER SECTION:
_ldap._tcp.server0.example.com. 86400 IN SRV 0 100 389 server0.example.com.

Received 154 bytes from 172.25.254.254#53 in 0 ms
```

Hosts and resource records

A typical host, whether a client or a server, will have the following records:

- One or more **A** and/or **AAAA** records mapping its host name to its IP addresses
- A **PTR** record for each of its IP addresses, reverse mapping them to its host name
- Optionally, one or more **CNAME** records mapping alternate names to its canonical host name

A DNS zone will typically have, in addition to the records for the hosts in the zone:

- Exactly one **SOA** record to specify how the zone works
- An **NS** record for each of its authoritative name servers
- One or more **MX** records mapping the domain name to the mail exchange which receives email for addresses ending in the domain name
- Optionally, one or more **TXT** records for functions such as SPF or Google Site Verification
- Optionally, one or more **SRV** records to locate services in the domain



References

host(1) and **dig(1)** man pages

RFC 1034: Domain names - concepts and facilities
<http://tools.ietf.org/html/rfc1034>

RFC 1035: Domain names - implementation and specification
<http://tools.ietf.org/html/rfc1035>

RFC 2181: Clarifications to the DNS Specification
<http://tools.ietf.org/html/rfc2181>

Practice: DNS Resource Record

Quiz

Match the following resource record types to the descriptions listed in the table.

A	AAAA	CNAME	MX	NS	PTR	SOA
SRV	TXT					

Purpose	Resource record type
Contains authoritative information for a zone, such as the email contact and several values that configure interactions between slave and master DNS servers.	
Maps host names to IPv4 addresses.	
Identifies the authoritative nameservers for a zone.	
Used to publish location of network services for a domain.	
Identifies the mail exchanges responsible for accepting emails for a domain.	
Maps host names to IPv6 addresses.	
Enables reverse DNS lookups of IP addresses to host names.	

Purpose	Resource record type
Aliases a name to a canonical name.	
Used to publish arbitrary human-readable text. Commonly used for SPF, DKIM, and DMARC.	

Solution

Match the following resource record types to the descriptions listed in the table.

Purpose	Resource record type
Contains authoritative information for a zone, such as the email contact and several values that configure interactions between slave and master DNS servers.	SOA
Maps host names to IPv4 addresses.	A
Identifies the authoritative nameservers for a zone.	NS
Used to publish location of network services for a domain.	SRV
Identifies the mail exchanges responsible for accepting emails for a domain.	MX
Maps host names to IPv6 addresses.	AAAA
Enables reverse DNS lookups of IP addresses to host names.	PTR
Aliases a name to a canonical name.	CNAME
Used to publish arbitrary human-readable text. Commonly used for SPF, DKIM, and DMARC.	TXT

Configuring a Caching Nameserver

Objectives

After completing this section, students should be able to configure a secure, caching nameserver using the **unbound** DNS server.

Caching nameservers and DNSSEC

Caching nameserver

Caching nameservers store DNS query results in a local cache and removes resource records from the cache when their TTLs expire. It is common to set up caching nameservers to perform queries on behalf of clients on the local network. This greatly improves the efficiency of DNS name resolutions by reducing DNS traffic across the Internet. As the cache grows, DNS performance improves as the caching nameserver answers more and more client queries from its local cache.

DNSSEC validation

Given the stateless nature of UDP, DNS transactions are prone to spoofing and tampering. Caching nameservers have historically been favored targets of attackers looking to redirect or hijack network traffic. This is often achieved by exploiting vulnerabilities in DNS server software to fool a DNS server into accepting and populating malicious data into its cache, a technique commonly referred to as *cache poisoning*. Once the attacker succeeds in poisoning a DNS server's cache, they effectively compromise the DNS data received by the numerous clients utilizing the caching name service on the DNS server and can consequently redirect or hijack the clients' network traffic.

While a caching nameserver can greatly improve DNS performance on the local network, they can also provide improved security by performing *Domain Name System Security Extensions (DNSSEC)* validation. DNSSEC validation enabled at the caching nameserver allows the authenticity and integrity of resource records to be validated prior to being placed in the cache for use by clients, and therefore protects clients against the consequences of cache poisoning.

Configuring and administering *unbound* as a caching nameserver

Several packages are available for configuring a caching nameserver, including *bind*, *dnsmasq*, and *unbound*. In this example, please follow along while the instructor demonstrates the configuration and administration of *unbound* as a secure, caching nameserver with DNSSEC validation enabled.

Configuring *unbound*

To configure **unbound** as a secure, caching nameserver:

1. Install **unbound**.

As *root*, install the **unbound** package.

```
[root@serverX ~]# yum install -y unbound
```

2. Start and enable **unbound.service**.

```
[root@serverX ~]# systemctl start unbound.service
[root@serverX ~]# systemctl enable unbound.service
ln -s '/usr/lib/systemd/system/unbound.service' '/etc/systemd/system/multi-
user.target.wants/unbound.service'
```

3. Configure the network interface to listen on.

By default, **unbound** only listens on the **localhost** network interface. To make **unbound** available to remote clients as a caching nameserver, use the **interface** option in the **server** clause of **/etc/unbound/unbound.conf** to specify the network interface(s) to listen on. A value of **0.0.0.0** will configure **unbound** to listen on all network interfaces:

```
interface: 0.0.0.0
```

4. Configure client access.

By default, **unbound** refuses recursive queries from all clients. In the **server** clause of **/etc/unbound/unbound.conf**, use the **access-control** option to specify which clients are allowed to make recursive queries.

```
access-control: 172.25.0.0/24 allow
```

5. Configure forwarding.

In **/etc/unbound/unbound.conf**, create a **forward-zone** clause to specify which DNS server(s) to forward queries to. DNS servers can be specified by host name using the **forward-host** option, or by IP address using the **forward-addr** option. For a caching nameserver, forward all queries by specifying a **forward-zone** of **."**.

```
forward-zone:
  name: "."
  forward-addr: 172.25.254.254
```

6. If desired, bypass DNSSEC validation for select unsigned zones.

By default, **unbound** is enabled to perform DNSSEC validation to verify all DNS responses received. The **domain-insecure** option in the **server** clause of **/etc/unbound/unbound.conf** can be used to specify a domain for which DNSSEC validation should be skipped. This is often desirable when dealing with an unsigned internal domain that would otherwise fail trust chain validation.

```
domain-insecure: example.com
```

7. If desired, install trust anchors for select signed zones without complete chain of trust.

Since not all ccTLDs have completed implementation of DNSSEC, the subdomains of these ccTLDs can be DNSSEC-signed but still have a broken *chain of trust*. This problem can be overcome by using the **trust-anchor** option in the **server** clause of **/etc/unbound/unbound.conf** to specify a trust anchor for the zone. Obtain the **DNSKEY**

record for the *key signing key (KSK)* of the zone using **dig** and input it as the value for the **trust-anchor** option.

```
[student@serverX ~]$ dig +dnssec DNSKEY example.com
```

```
trust-anchor: "example.com. 3600 IN DNSKEY 257 3 8 AwEAAawt7Hp1I5M8GGAsxuyCyjF01
+QlcgVN11CRZ4vP66qbDCX0BnSh Z11BGb//4zSG/8mmBHirL2FLg+mVuIixig
+iroZYjh4iTKV0hv2hZftR wyrQHK++qXvCCWN3ki51RG/e8R4k0EV71rZ80gQvPWx6F91qroq0Ppcf
7PPxippeH0n+PxnP0hpyLyo1mx1rPs/cMpL3j0MufGP+LJYh+fBU7lt0
sP5i09HaJPruzyZML9Btpv8ZAdQhwtXVG0+MnET2qT/1+TljpxZn6ye egFCRFRHBJmo6iirJnUwra/
k1krgEn2Q+BXGTOMTTKQdYz40xYEa1z7a pu3a09dYnBM="
```

8. Save changes to **/etc/unbound/unbound.conf**.
9. Check the **/etc/unbound/unbound.conf** configuration file for syntax errors.

```
[root@serverX ~]# unbound-checkconf
unbound-checkconf: no errors in /etc/unbound/unbound.conf
```

10. Restart **unbound.service**.

```
[root@serverX ~]# systemctl restart unbound.service
```

11. Configure the firewall to allow DNS traffic.

```
[root@serverX ~]# firewall-cmd --permanent --add-service=dns
success
[root@serverX ~]# firewall-cmd --reload
success
```

Dumping and loading unbound cache

Administrators of caching nameservers need to dump out cache data when troubleshooting DNS issues, such as those resulting from stale resource records. With an **unbound** DNS server, the cache can be dumped by running the **unbound-control** utility in conjunction with the **dump_cache** subcommand.

```
[root@serverX ~]# unbound-control dump_cache
START_RRSET_CACHE
;rrset 86395 1 0 3 3
classroom.example.com. 86395 IN A 172.25.254.254
;rrset 86395 1 0 7 3
example.com. 86395 IN NS classroom.example.com.
;rrset 86395 1 0 8 3
example.com. 86395 IN A 172.25.254.254
END_RRSET_CACHE
START_MSG_CACHE
msg example.com. IN A 33152 1 86395 3 1 1 1
example.com. IN A 0
example.com. IN NS 0
classroom.example.com. IN A 0
END_MSG_CACHE
EOF
```

Executing **unbound-control** with the **dump_cache** command dumps out the cache to *stdout* in a text format. This output can be directed to a file for storage and be loaded back into cache later with **unbound-control load_cache**, if desired. **unbound-control load_cache** reads from *stdin* to populate the cache.

```
[root@serverX ~]# unbound-control load_cache < dump.out
ok
```

Flushing unbound cache

Administrators of caching nameservers also need to purge outdated resource records from cache from time to time. Erroneous and outdated resource records in cache will keep their newly corrected counterparts from becoming available to clients until the TTLs on the outdated resource records expire. Rather than waiting for TTL expiration, administrators can forcibly purge the outdated records from cache by executing **unbound-control** with the **flush** subcommand.

```
[root@serverX ~]# unbound-control flush www.example.com
ok
```

If all resource records belonging to a domain need to be purged from the cache of an **unbound** DNS server, **unbound-control** can be executed with the **flush_zone** subcommand.

```
[root@serverX ~]# unbound-control flush_zone example.com
ok removed 3 rrssets, 1 messages and 0 key entries
```

Updating Local Caching unbound Configuration with dnsssec-trigger

In addition to providing caching name service for a local subnet, **unbound** can also be useful as a local caching nameserver to provide secure DNS name resolution for local use on an individual system. For a local caching nameserver setup, the **nameserver** entry in **/etc/resolv.conf** will be configured to point to localhost where **unbound** is listening. The **unbound** configuration will forward DNS requests to upstream nameservers and validate their responses.

For DHCP systems running local caching name service, the upstream nameservers specified in **unbound**'s configuration may become outdated if DNS servers provided by DHCP change. The **dnsssec-trigger** tool supplied by the package of the same name can be leveraged to automatically update forwarder settings in **unbound**'s configuration file to point to the new DNS servers. The use of the **dnsssec-trigger** tool in conjunction with **unbound** is mostly useful for secure DNS name resolution on roaming client machines.

R

References

unbound(8), **unbound-checkconf**(8), **unbound.conf**(5), **unbound-control**(8), **dnsssec-trigger**(8) man pages

Practice: Configuring unbound as a Caching Nameserver

Guided exercise

In this lab, you will configure **unbound** as a caching nameserver and administer its cache data.

Resources:	
Files:	/etc/unbound/unbound.conf
Machines:	<ul style="list-style-type: none"> • desktopX • serverX

Outcomes:

The **unbound** service enabled and configured as a caching nameserver on **serverX** to listen on interface **172.25.X.11** and to accept connections from **172.25.X.0/24**. Recursive queries are forwarded to **172.25.254.254**. The **example.com** zone is exempted from DNSSEC validation.

Before you begin...

- Reset the **desktopX** system.
- Reset the **serverX** system.
- Log into the **desktopX** system.
- Log into the **serverX** system and switch to root using **sudo -i**.

You have been asked to improve name resolution performance and enhance DNS security in your company's datacenter. You have decided to deploy a caching nameserver using an **unbound** DNS server. You will configure **unbound** to respond only to queries on the datacenter subnet, **172.25.X.0/24**. Recursive queries will be forwarded to the company's main internal DNS server at **172.25.254.254**. This main internal DNS server hosts an internal, split DNS copy of the company's **example.com** zone. This **example.com** zone is not DNSSEC-signed, so it will need to be exempted from DNSSEC validation on your caching nameserver. After configuration of the caching nameserver is complete on **serverX**, you will test it by querying for the host names **desktopX.example.com** and **serverX.example.com** from **desktopX**. Verify that the queries made are populated into the nameserver's cache. Lastly, you will purge **serverX.example.com** from the cache.

- 1. Install the *unbound* package on **serverX**.

```
[root@serverX ~]# yum install -y unbound
```

- 2. Enable and start the **unbound** service.

```
[root@serverX ~]# systemctl enable unbound.service
ln -s '/usr/lib/systemd/system/unbound.service' '/etc/systemd/system/multi-user.target.wants/unbound.service'
[root@serverX ~]# systemctl start unbound.service
```

- 3. Configure **unbound** to allow queries from the **172.25.X.0/24** subnet, exempt the **example.com** zone from DNSSEC validation, and forward all queries to **172.25.254.254**.
 - 3.1. By default, **unbound** listens on the loopback interface. Configure **unbound** to listen on the **172.25.X.11** interface on **serverX** by adding the following option in the **server** clause of **/etc/unbound/unbound.conf**.

```
interface: 172.25.X.11
```

- 3.2. Allow queries from the **172.25.X.0/24** subnet by adding the following option in the **server** clause of **/etc/unbound/unbound.conf**.

```
access-control: 172.25.X.0/24 allow
```

- 3.3. Exempt the **example.com** zone from DNSSEC validation by adding the following option in the **server** clause of **/etc/unbound/unbound.conf**.

```
domain-insecure: "example.com"
```

- 3.4. Forward all queries to **172.25.254.254** by adding a **forward-zone** clause to the end of the **/etc/unbound/unbound.conf** file.

```
forward-zone:  
    name: .  
    forward-addr: 172.25.254.254
```

- 4. Check **/etc/unbound/unbound.conf** for syntax errors.

```
[root@serverX ~]# unbound-checkconf  
unbound-checkconf: no errors in /etc/unbound/unbound.conf
```

- 5. Restart the **unbound** service.

```
[root@serverX ~]# systemctl restart unbound.service
```

- 6. Configure the firewall to allow DNS traffic.

```
[root@serverX ~]# firewall-cmd --permanent --add-service=dns  
success  
[root@serverX ~]# firewall-cmd --reload  
success
```

- 7. Verify the caching name service by performing queries and examining the contents of the cache.

- 7.1. Dump the cache to see its contents.

```
[root@serverX ~]# unbound-control dump_cache
START_RRSET_CACHE
END_RRSET_CACHE
START_MSG_CACHE
END_MSG_CACHE
EOF
```

- 7.2. From **desktopX**, query **serverX** for the **A** record of host name **desktopX.example.com**.

```
[student@desktopX ~]$ dig @serverX.example.com A desktopX.example.com
...
desktopX.example.com. 86349 IN A 172.25.X.10
...
```

- 7.3. From **desktopX**, query **serverX** for the **A** record of host name **serverX.example.com**.

```
[student@desktopX ~]$ dig @serverX.example.com A serverX.example.com
...
serverX.example.com. 86364 IN A 172.25.X.11
...
```

- 7.4. On **serverX**, dump out the cache again. You should see the queried records in the cache.

```
[root@serverX ~]# unbound-control dump_cache
START_RRSET_CACHE
;rrset 85886 1 0 8 3
serverX.example.com. 85886 IN A 172.25.X.11
;rrset 85878 1 0 8 X
desktopX.example.com. 85878 IN A 172.25.X.10
;rrset 85878 1 0 7 3
example.com. 85878 IN NS classroom.example.com.
;rrset 85878 1 0 3 3
classroom.example.com. 85878 IN A 172.25.254.254
END_RRSET_CACHE
START_MSG_CACHE
msg desktopX.example.com. IN A 33152 1 85878 3 1 1 1
desktopX.example.com. IN A 0
example.com. IN NS 0
classroom.example.com. IN A 0
msg serverX.example.com. IN A 33152 1 85886 3 1 1 1
serverX.example.com. IN A 0
example.com. IN NS 0
classroom.example.com. IN A 0
END_MSG_CACHE
EOF
```

- 7.5. Purge the **serverX.example.com** record from the cache.

```
[root@serverX ~]# unbound-control flush serverX.example.com
ok
```


- 7.6. On **serverX**, dump out the cache again. You should no longer see the **A** record for **serverX.example.com** in the cache.

```
[root@serverX ~]# unbound-control dump_cache
START_RRSET_CACHE
;rrset 85878 1 0 8 X
desktopX.example.com. 85878 IN A 172.25.X.10
;rrset 85878 1 0 7 3
example.com. 85878 IN NS classroom.example.com.
;rrset 85878 1 0 3 3
classroom.example.com. 85878 IN A 172.25.254.254
END_RRSET_CACHE
START_MSG_CACHE
msg desktopX.example.com. IN A 33152 1 85878 3 1 1 1
desktopX.example.com. IN A 0
example.com. IN NS 0
classroom.example.com. IN A 0
msg serverX.example.com. IN A 33152 1 85886 3 1 1 1
serverX.example.com. IN A 0
example.com. IN NS 0
classroom.example.com. IN A 0
END_MSG_CACHE
EOF
```

DNS Troubleshooting

Objectives

After completing this section, students should be able to:

- Use **dig** to troubleshoot common DNS problems.
- Identify symptoms and causes associated with common DNS issues.

Troubleshooting DNS

Due to the client-server architecture of DNS, properly working DNS name resolution on a system is almost always dependent on not only the proper configuration and operation of DNS on that system, but also on that of its resolving nameserver and the many authoritative nameservers used to resolve its DNS requests. Since DNS is a distributed directory, recursive name resolution often involves numerous behind-the-scenes interactions with many different authoritative nameservers. These numerous interactions create many possible points for failure.

The use of caching nameservers significantly reduces DNS workloads and improves DNS performance. However, the caching function adds another point of failure by creating scenarios where it is possible for DNS responses received by clients to be inaccurate due to the data no longer being current.

Due to the critical role that DNS plays in the functioning of network services, it is important that Linux administrators be able to quickly resolve DNS issues when they occur, in order to minimize service interruptions. The key to accurate and efficient DNS troubleshooting is being able to pinpoint which of the multiple points in the myriad of behind-the-scenes client-server interactions is responsible for the unexpected behavior observed. This requires the use of proper tools and a clear understanding of the diagnostic data they provide. *Domain Internet Groper (dig)* is a good tool for investigating DNS issues due to its verbose diagnostic output.

Name resolution methods

Because DNS service is often the most widely used method of name resolution, it often bears the blame whenever unexpected name resolution results occur. Remember that aside from DNS, in a heterogeneous environment, name resolution on networked hosts can occur via other methods, such as local hosts files, Windows Internet Name Service (WINS), etc.

On Linux systems, name resolution is attempted first with the hosts file **/etc/hosts** by default, per order specified in **/etc/nsswitch.conf**. Therefore, when beginning name resolution troubleshooting, do not leap to the assumption that the issue resides with DNS. Begin first by identifying the name resolution mechanism which is in play, rather than simply starting with DNS. The **getent** command from the *glibc-common* package, as well as the **gethostip** command from the *syslinux* package, can both be used to perform name resolution, mirroring the process used by most applications in following the order of host name resolution as dictated by **/etc/nsswitch.conf**.

```
[student@desktopX ~]$ getent hosts example.com
172.25.254.254 example.com
```

```
[student@desktopX ~]$ gethostip example.com
example.com 172.25.254.254 AC19FEFE
```

If the result of **getent** or **gethostip** differs from that produced by **dig**, then it's a clear indication that something besides DNS is responsible for the unexpected name resolution result. Consult **/etc/nsswitch.conf** to determine what other name resolution mechanisms are employed before DNS.

Client-server network connectivity

For DNS name resolution to work properly, a system must first be able to conduct client-server interactions with its resolving nameserver or other authoritative nameservers. Some common DNS issues that have their origin at this layer are the result of resolver and firewall misconfigurations.

When using **dig** to troubleshoot a DNS issue, if a response is not received from a DNS server, it is a clear indication that the cause lies with the client-server network connectivity to the DNS server.

```
[student@desktopX ~]$ dig A example.com
; <<> DiG 9.9.4-RedHat-9.9.4-14.el7 <<> A example.com
;; global options: +cmd
;; connection timed out; no servers could be reached
```

A possible cause is the inability to reach the DNS server due to incorrect DNS server IP address(es) in a system's DNS configuration. This could be in **/etc/resolv.conf** in the case of a system acting as a DNS client or in the **forward-zone** clause of **/etc/unbound/unbound.conf** in the case of a system configured as an **unbound** caching nameserver.

Another possible cause are firewall rules, on either the client or server system, blocking DNS traffic on port 53. While DNS mostly uses the UDP protocol, it is important to note that when response data sizes exceed 512 bytes, or 4096 bytes in the case of DNS servers that support *Extension Mechanism for DNS (EDNS)*, resolvers will fall back to using TCP to retry the query. Therefore, proper DNS configuration should allow for DNS traffic on port 53 for both UDP and TCP. Allowing port 53 traffic for UDP only will result in a truncation error when the resolver encounters a response that is larger than what it can handle over UDP.

```
[student@desktopX ~]$ dig @serverX.example.com A labhost1.example.com
;; Truncated, retrying in TCP mode.
;; Connection to 172.25.1.11#53(172.25.1.11) for labhost1.example.com failed:
host unreachable.
```

dig's **tcp** or **vc** options are helpful for troubleshooting whether DNS queries can succeed with TCP. These options force **dig** to use TCP, rather than the default behavior of using UDP first and falling back to TCP only when response size necessitates it.

```
[student@desktopX ~]$ dig +tcp A example.com
```

When dealing with DNS issues at the network layer, **dig** provides very sparse output and it is therefore often useful to also use a network packet analyzer, such as **tcpdump**, to determine what is transpiring behind the scenes at the network layer. Using **tcpdump**, the administrator can determine information that cannot be ascertained with **dig** alone, such as the destination IP address of the DNS request, if request packets leave the client, if request packets reach the server, if response packets leave the server, if response packets reach the client, etc.

DNS response codes

If DNS client-server communication is successful, **dig** will generate much more verbose output detailing the nature of the response received from the DNS server. The **status** field in the **HEADER** section of **dig**'s output reports the response code generated by the DNS server in response to the client's query.

```
[student@desktopX ~]$ dig A example.com
; <<>> DiG 9.9.4-RedHat-9.9.4-14.el7 <<>> A example.com
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 30523
...
```

A status of **NOERROR** indicates the query was resolved successfully. If the server encounters problems fulfilling the client's query, one of the following common error statuses will be displayed.

DNS Return Codes

Code	Meaning
SERVFAIL	The nameserver encountered a problem while processing the query.
NXDOMAIN	The queried name does not exist in the zone.
REFUSED	The nameserver refused the client's DNS request due to policy restrictions.

SERVFAIL

A common cause of **SERVFAIL** status is the failure of the DNS server to communicate with the nameservers authoritative for the name being queried. This may be due to the authoritative nameservers being unavailable. It could also be a problem at the network layer interfering with the client-server communication between the DNS server and the authoritative nameservers, such as network routing issues or firewall rules at any hop in the network path.

To determine why a nameserver is generating a **SERVFAIL** status while performing recursion on behalf of a client's query, the administrator of the nameserver will need to determine which nameserver communication is causing the failure. **dig**'s **+trace** option is helpful in this scenario to see the results of nameserver's iterative queries starting with the root nameservers.

NXDOMAIN

An **NXDOMAIN** status indicates that no records were found associated with the name queried. If this is not the expected result and the query is directed at a server that is not authoritative for the name, then the server's cache may contain a negative cache for the name. The user can either wait for the server to expire the negative cache of that name, or submit a request to the server administrator to flush the name from the server's cache. Once the name is removed from cache, the server will query the authoritative nameserver to receive current resource records for the name.

The other scenario where an **NXDOMAIN** status may be unexpectedly encountered is when querying a **CNAME** record containing an *orphaned CNAME*. In a **CNAME** record, the name on the right side of the record, the canonical name, should point to a name that contains either **A** or **AAAA** records. If these associated **A** or **AAAA** records are nonexistent or are later removed, then the canonical name in the **CNAME** record is orphaned. When this occurs, queries for the owner

name in the **CNAME** record will no longer be resolvable and will result in an **NXDOMAIN** return code.

REFUSED

A **REFUSED** status indicates that the DNS server has a policy restriction which keeps it from fulfilling the client's query. Policy restrictions are often implemented on DNS servers to restrict which clients can make recursive queries and zone transfer requests. Some common causes of an unexpected **REFUSED** return code are clients configured to query the wrong DNS servers or DNS server misconfiguration causing valid client requests to be refused.

Other common DNS issues

Outdated cached data

A DNS return code of **NOERROR** signifies that no errors were encountered in resolving a query. However, it does not guarantee that there are no DNS issues present. There are situations where the DNS records in the DNS response may not match the expected result. The most common cause for an incorrect answer is that the answer originated from cached data, which is no longer current.

In these situations, first confirm that the response is indeed nonauthoritative cached data. This can be easily determined by looking at the **flags** section of **dig**'s output. DNS responses, which are authoritative answers, will be indicated by the presence of the **aa** flag.

```
[student@desktopX ~]$ dig A example.com
; <<> DiG 9.9.4-RedHat-9.9.4-14.e17 <<> A example.com
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 22257
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 2
```

DNS responses that originate from cached data are not authoritative and therefore will not have the **aa** flag set. The other telltale sign that an answer is coming from cache is the counting down of the resource record's TTL value in the responses of each subsequent query. TTLs of cached data will continuously count down to expiration while TTLs of authoritative data will always remain static.

Responses for nonexistent records

If a record has been removed from a zone and a response is still received when querying for the record, first confirm that the queries are not being answered from cached data. If the answer is authoritative as confirmed by the presence of the **aa** flag in **dig**'s output, then a possible cause is the presence of a *wildcard* (*) record in the zone.

```
*.example.com. IN A 172.25.254.254
```

A wildcard record serves as a catchall for all queries of a given type for a nonexistent name. With the previous wildcard record in place, if an **A** record previously existed for **serverX.example.com** and it is removed, queries for the name will still succeed and the IP address in the wildcard **A** record will be provided in its place.

Non-FQDN name errors

In a zone file, host names which are not expressed as *Fully Qualified Domain Names (FQDNs)* are automatically expanded to FQDNs by appending the name of the zone. To indicate that a name

is an FQDN in a zone file, it must be ended with a ".", i.e., "**www.example.com.**". Failure to do so can lead to different issues depending on the type of record that the mistake is made in. For example, such a mistake made in the type-specific data portion of **NS** records have the potential of incapacitating an entire zone, while a mistake made in **MX** records could cause a complete halt of email delivery for a domain.

Looping **CNAME** records

While technically feasible, **CNAME** records that point to **CNAME** records should be avoided to reduce DNS lookup inefficiency. Another reason this is undesirable is the possibility of creating unresolvable **CNAME** loops, such as:

```
test.example.com. IN CNAME lab.example.com.  
lab.example.com. IN CNAME test.example.com.
```

While a **CNAME** record with an orphaned **CNAME** will result in an **NXDOMAIN** status, looping **CNAME** records will return as **NOERROR**.

Missing **PTR** records

Many network services use DNS to perform reverse lookups of incoming connections from clients. The absence of **PTR** records in DNS may result in issues, the nature of which varies depending on the service. SSHD, by default, will perform reverse lookups of connecting client IPs. Absence of **PTR** records will lead to delays in the establishment of these connections.

Many MTAs incorporate reverse DNS lookups of connecting client IPs as a defense against malicious email clients. In fact, many MTAs are configured to reject client connections for IPs, which cannot be resolved with a **PTR** query in DNS. As such, administrators supporting network services need to ensure that they understand the requirements these services have for not just forward, but also reverse DNS lookups.

Round-robin DNS

A name can have multiple **A** or **AAAA** records in DNS. This is known as *round-robin DNS*, and is often used as a simple, low-cost, load-balancing mechanism to distribute network resource loads across multiple hosts. When a DNS client queries for a name that contains multiple **A** or **AAAA** records, all records are returned as a set. However, the order that the records are returned in the set permutes for each query. Since clients normally make use of the first address in the set, the variation in the order of the records in each response effectively results in a distribution of network service requests across the multiple IP addresses in these round-robin DNS records.

While round-robin DNS is a valid technical configuration, there are times when this configuration is inadvertently created. When a request to change the IP address of an **A** record is mistakenly implemented as a resource record addition rather than a resource record modification, then round-robin DNS is created. If the network resources on the old IP address is retired, the load distribution effect of the round-robin DNS will result in service failures for approximately half of the clients.

R

References

dig(1), **getent**(1), **gethostip**(1) man pages

Practice: Troubleshooting DNS

Guided exercise

In this lab, you will troubleshoot and resolve a name resolution issue by systematically verifying name service configurations to pinpoint the cause.

Resources:	
Files:	<ul style="list-style-type: none"> • /etc/nsswitch.conf • /etc/hosts • /etc/resolv.conf
Machines:	<ul style="list-style-type: none"> • desktopX

Outcomes:

The root cause for failed name resolution for **example.com** on **desktopX** will be identified and fixed. After the issue is resolved, name resolution for **example.com** on **desktopX** will succeed.

Before you begin...

- Reset the **desktopX** system.
- Log into and set up the **desktopX** system.

```
[student@desktopX ~]$ lab dnsissue2 setup
```

A user reports that there is an issue occurring when an SSH session is initiated to **example.com** from **desktopX** and a "**Could not resolve hostname example.com: Name or service not known**" error is being generated. Therefore, the user complains that there is a problem with DNS name resolution of **example.com**. Like all other hosts on the network, **desktopX** should be using **172.25.254.254** for DNS resolution. You will troubleshoot the issue, identify the root cause, apply a fix, and then verify that the problem is resolved.



Note

You can not actually login to **example.com** using SSH. Seeing a login prompt is all that is required, and the login prompt can be cancelled by pressing **Ctrl+C**

- 1. Replicate the reported issue by attempting an SSH session to **example.com** from **desktopX**.

```
[student@desktopX ~]$ ssh example.com
ssh: Could not resolve hostname example.com: Name or service not known
```

- 2. Verify the result of name resolution for **example.com**.

```
[student@desktopX ~]$ getent hosts example.com
[student@desktopX ~]$
```

- 3. Verify configurations which affect name resolution.

- 3.1. Verify the order that name services are used.

```
[student@desktopX ~]$ grep ^hosts: /etc/nsswitch.conf
hosts:      files dns
```

- 3.2. Since **files** are used first, verify the contents of **/etc/hosts**.

```
[student@desktopX ~]$ grep [[:space:]]example.com /etc/hosts
[student@desktopX ~]$
```

- 3.3. Since no hosts file entry exist for **example.com**, verify the contents of **/etc/resolv.conf**; you should see that an incorrect nameserver IP is the cause of the name resolution failure.

```
[student@desktopX ~]$ grep ^nameserver /etc/resolv.conf
nameserver 172.25.254.255
```

```
[student@desktopX ~]$ dig @172.25.254.255 A example.com
; <<>> DiG 9.9.4-RedHat-9.9.4-14.el7 <<>> @172.25.254.255 A example.com
; (1 server found)
;; global options: +cmd
;; connection timed out; no servers could be reached
```

- 4. As **root**, fix the error in **/etc/resolv.conf** and verify that this resolves the name resolution issue.

- 4.1. Since the **nameserver** entry populated by DHCP appears to have been manually modified, as **root**, force a refresh of the data from DHCP and validate that the entry is fixed.

```
[student@desktopX ~]$ sudo systemctl restart NetworkManager
```

```
[student@desktopX ~]$ grep ^nameserver /etc/resolv.conf
nameserver 172.25.254.254
```

- 4.2. Verify the results of name resolution for **example.com**.

```
[student@desktopX ~]$ getent hosts example.com
172.25.254.254 example.com
```



```
[student@desktopX ~]$ dig A example.com
; <<> DiG 9.9.4-RedHat-9.9.4-14.el7 <<> example.com
;; global options: +cmd
;; Got answer:
;; -->HEADER<<-- opcode: QUERY, status: NOERROR, id: 36048
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 2

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
;example.com.    IN A

;; ANSWER SECTION:
example.com.    86400 IN A 172.25.254.254

;; AUTHORITY SECTION:
example.com.    86400 IN NS classroom.example.com.

;; ADDITIONAL SECTION:
classroom.example.com. 86400 IN A 172.25.254.254

;; Query time: 1 msec
;; SERVER: 172.25.254.254#53(172.25.254.254)
;; WHEN: Thu May 15 07:28:35 EDT 2014
;; MSG SIZE rcvd: 96
```

- 4.3. Verify that SSH connection to **example.com** from **desktopX** now succeeds.

```
[student@desktopX ~]$ ssh example.com
The authenticity of host 'example.com (172.25.254.254)' can't be
established.
ECDSA key fingerprint is 12:b3:c8:3e:6b:d2:9f:43:67:a5:f2:2a:f0:7c:2f:b6.
Are you sure you want to continue connecting (yes/no)?
```

Lab: Managing DNS for Servers

Performance checklist

In this lab, you will use DNS troubleshooting techniques and knowledge of **unbound** configuration to resolve DNS issues with a misconfigured **unbound** caching nameserver.

Resources:	
Files:	• /etc/unbound/ubound.conf
Machines:	• desktopX • serverX

Outcomes:

The resolution of issues with DNS requests performed from **desktopX** against the **unbound** caching nameserver running on **serverX**. Properly configured **unbound** caching nameserver running on **serverX**.

Before you begin...

- Reset your **desktopX** system.
- Reset your **serverX** system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab dnsunboundissues setup
```

To improve DNS performance and security, a fellow system administrator has recently configured a secure caching nameserver running the **unbound** DNS server on **serverX** to serve your local subnet, **172.25.X.0/24**. The caching nameserver uses DNS server **172.25.254.254** for name resolution. The administrator released the caching nameserver for use just prior to leaving for a week-long vacation.

Soon after, you receive reports of DNS issues from users trying to use the newly configured caching nameserver for the first time. One user provides a specific example and reports that name resolution for **example.com** fails.

You will conduct troubleshooting to determine and correct the **unbound** misconfiguration that is causing this issue.

1. On **desktopX**, replicate the reported issue to determine the nature and scope of the problem.
2. Verify the operation of the caching nameserver on **serverX** by issuing queries to it locally from **serverX**.
3. Fix the issue discovered and verify that it resolved the name resolution issue.
4. While the DNS communication now works, the return code in the DNS response indicates that there may be another misconfiguration issue present. Verify the **unbound** configuration on **serverX**.

5. Fix the issue discovered and verify that it resolved the name resolution issue.
6. On **desktopX**, run **lab dnsunboundissues grade** to verify the **unbound** caching nameserver on **serverX** is working properly.

Solution

In this lab, you will use DNS troubleshooting techniques and knowledge of **unbound** configuration to resolve DNS issues with a misconfigured **unbound** caching nameserver.

Resources:	
Files:	• /etc/unbound/unbound.conf
Machines:	• desktopX • serverX

Outcomes:

The resolution of issues with DNS requests performed from **desktopX** against the **unbound** caching nameserver running on **serverX**. Properly configured **unbound** caching nameserver running on **serverX**.

Before you begin...

- Reset your **desktopX** system.
- Reset your **serverX** system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab dnsunboundissues setup
```

To improve DNS performance and security, a fellow system administrator has recently configured a secure caching nameserver running the **unbound** DNS server on **serverX** to serve your local subnet, **172.25.X.0/24**. The caching nameserver uses DNS server **172.25.254.254** for name resolution. The administrator released the caching nameserver for use just prior to leaving for a week-long vacation.

Soon after, you receive reports of DNS issues from users trying to use the newly configured caching nameserver for the first time. One user provides a specific example and reports that name resolution for **example.com** fails.

You will conduct troubleshooting to determine and correct the **unbound** misconfiguration that is causing this issue.

1. On **desktopX**, replicate the reported issue to determine the nature and scope of the problem.
 - 1.1. From **desktopX**, issue a query for the address of **example.com** to the caching nameserver on **serverX**.

```
[student@desktopX ~]$ dig @serverX.example.com A example.com
; <<>> DiG 9.9.4-RedHat-9.9.4-14.e17 <<>> @serverX.example.com example.com
; (1 server found)
;; global options: +cmd
;; connection timed out; no servers could be reached
```

2. Verify the operation of the caching nameserver on **serverX** by issuing queries to it locally from **serverX**.

- 2.1. Issue the query to the **localhost** interface. The query succeeds with no errors.

```
[student@serverX ~]$ dig @localhost A example.com
; <<>> DiG 9.9.4-RedHat-9.9.4-14.el7 <<>> @localhost A example.com
; (2 servers found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 64095
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
;example.com.      IN A

;; ANSWER SECTION:
example.com.      84526 IN A 172.25.254.254

;; Query time: 0 msec
;; SERVER: ::1#53(::1)
;; WHEN: Wed May 21 05:12:00 EDT 2014
;; MSG SIZE rcvd: 56
```

- 2.2. Issue the query to the **172.25.X.11** interface. No response is received and the query times out.

```
[student@serverX ~]$ dig @172.25.X.11 A example.com
; <<>> DiG 9.9.4-RedHat-9.9.4-14.el7 <<>> @172.25.X.11 A example.com
; (1 server found)
;; global options: +cmd
;; connection timed out; no servers could be reached
```

- 2.3. Determine which interfaces **unbound** is listening on. You should discover that it's only listening on the **localhost** interface.

```
[student@serverX ~]$ sudo ss -tulpn | grep -w 53
tcp    UNCONN    0      0      127.0.0.1:53      *:*
users:(("unbound",2192,5))
tcp    UNCONN    0      0      :::1:53          :::*
users:(("unbound",2192,3))
tcp    LISTEN    0      5      127.0.0.1:53      *:*
users:(("unbound",2192,6))
tcp    LISTEN    0      5      :::1:53          :::*
users:(("unbound",2192,4))
```

3. Fix the issue discovered and verify that it resolved the name resolution issue.
- 3.1. Configure **unbound** to listen on all interfaces by adding the following entry in **/etc/unbound/unbound.conf**.

```
interface: 0.0.0.0
```

- 3.2. Check the configuration for syntax errors.

```
[student@serverX ~]$ sudo unbound-checkconf
unbound-checkconf: no errors in /etc/unbound/unbound.conf
```

- 3.3. Restart the service for the changes to take place.

```
[student@serverX ~]$ sudo systemctl restart unbound
```

- 3.4. On **serverX**, rerun the query to the **172.25.X.11** interface. You should now receive a response, but with a status of **REFUSED**.

```
[student@serverX ~]$ dig @172.25.X.11 A example.com
; <<> DiG 9.9.4-RedHat-9.9.4-14.el7 <<> @172.25.X.11 A example.com
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->HEADER<<- opcode: QUERY, status: REFUSED, id: 50719
;; flags: qr rd ad; QUERY: 0, ANSWER: 0, AUTHORITY: 0, ADDITIONAL: 0
;; WARNING: recursion requested but not available

;; Query time: 0 msec
;; SERVER: 172.25.X.11#53(172.25.X.11)
;; WHEN: Wed May 21 05:36:43 EDT 2014
;; MSG SIZE rcvd: 12
```

4. While the DNS communication now works, the return code in the DNS response indicates that there may be another misconfiguration issue present. Verify the **unbound** configuration on **serverX**.
- 4.1. Since **REFUSED** status indicates a policy restriction, verify that the proper access policy is in place in **/etc/unbound/unbound.conf**. You should see that no access control has been granted for the subnet.

```
[student@serverX ~]$ sudo grep ^[:,space:]*access-control /etc/unbound/
unbound.conf
```

5. Fix the issue discovered and verify that it resolved the name resolution issue.
- 5.1. Grant the **172.25.X.0/24** subnet access to the caching nameserver in **/etc/unbound/unbound.conf**.

```
[student@serverX ~]$ sudo grep ^[:,space:]*access-control /etc/unbound/
unbound.conf
access-control: 172.25.X.0/24 allow
```

- 5.2. Check the configuration for syntax errors.

```
[student@serverX ~]$ sudo unbound-checkconf
unbound-checkconf: no errors in /etc/unbound/unbound.conf
```

- 5.3. Restart the service for the changes to take place.

```
[student@serverX ~]$ sudo systemctl restart unbound
```

5.4. On **serverX**, rerun the previously failed query. It should now succeed.

```
[student@serverX ~]$ dig @172.25.X.11 A example.com
; <<>> DiG 9.9.4-RedHat-9.9.4-14.el7 <<>> @172.25.X.11 A example.com
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 25229
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
;example.com.      IN A

;; ANSWER SECTION:
example.com.      86339 IN A 172.25.254.254

;; Query time: 0 msec
;; SERVER: 172.25.X.11#53(172.25.X.11)
;; WHEN: Wed May 21 05:56:16 EDT 2014
;; MSG SIZE rcvd: 56
```

5.5. On **desktopX**, rerun the query to **serverX**. It should now succeed.

```
[student@desktopX ~]$ dig @serverX.example.com A example.com
; <<>> DiG 9.9.4-RedHat-9.9.4-14.el7 <<>> @serverX.example.com A example.com
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 25229
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
;example.com.      IN A

;; ANSWER SECTION:
example.com.      86339 IN A 172.25.254.254

;; Query time: 0 msec
;; SERVER: 172.25.X.11#53(172.25.X.11)
;; WHEN: Wed May 21 05:56:16 EDT 2014
;; MSG SIZE rcvd: 56
```

6. On **desktopX**, run **lab dnsunboundissues grade** to verify the **unbound** caching nameserver on **serverX** is working properly.

```
6.1. [student@desktopX ~]$ lab dnsunboundissues grade
PASS
```

Summary

DNS Concepts

In this section, students learned how to:

- Review the structure of the Domain Name System.
- Detail the anatomy of DNS lookups.
- Identify common DNS resource records and their uses.

Configuring a Caching Nameserver

In this section, students learned how to:

- Identify the need for DNSSEC validation on a caching nameserver.
- Configure **unbound** DNS server as a caching nameserver.
- Administer **unbound**'s cache data.

DNS Troubleshooting

In this section, students learned how to:

- Identify symptoms and causes of common DNS issues.
- Interpret results of diagnostic tools used in DNS troubleshooting.
- Define common DNS response codes.



CHAPTER 6

CONFIGURING EMAIL TRANSMISSION

Overview	
Goal	To relay all email sent by the system through an SMTP gateway.
Objectives	<ul style="list-style-type: none">• Configure a Red Hat Enterprise Linux server to transmit all email through an unauthenticated SMTP gateway.
Sections	<ul style="list-style-type: none">• Configuring a Send-only Email Configuration (and Practice)
Lab	<ul style="list-style-type: none">• Configuring Email Transmission

Configuring Send-only Email Service

Objectives

After completing this section, students should be able to:

Configure a Red Hat Enterprise Linux server to transmit all email through an unauthenticated SMTP gateway.

Email architecture and null clients

In today's corporate environments, email is a common method of communication. End users may use dedicated *mail clients* such as **Evolution** and **mutt** to read and send email, or the organization may have a web-based interface for its mail service.

However, Linux servers also send email, usually for automatic purposes or to report errors to an administrator. They generally use, directly or indirectly, a standard program called **/usr/sbin/sendmail** (provided in RHEL 7 by **Postfix**) to send these messages.

In practice, most servers are monitored and send out mails when incidents occur. This often requires a configured **/usr/sbin/sendmail** to send emails to notify the responsible system administrators by using the corporate SMTP server to transmit the messages.

A *null client* is a client machine that runs a local mail server which forwards all emails to an outbound mail relay for delivery. A null client does not accept local delivery for *any* messages, it can only send them to the outbound mail relay. Users may run mail clients on the null client to read and send emails.

This section will look at how to configure a RHEL 7 server as a Postfix null client, which will use **sendmail** and the SMTP protocol to transmit mail messages to the outside world through an existing outgoing mail server.

Transmission of an email message

To send an email, in most cases the mail client communicates with an outgoing mail server, which will help *relay* that message to its final destination. The mail client transmits messages to the mail server using the Simple Mail Transfer Protocol (SMTP).

The outgoing mail relay may require no authentication from internal clients, in which case the server listens on port 25/TCP. In that case, the relay will restrict which hosts can relay through IP address based restrictions or firewall rules.



Note

In cases where the outbound SMTP relay is reachable from the Internet, it is normally configured as a *mail submission agent* (MSA) for security and anti-spam reasons. An MSA listens on port 587/TCP and requires authentication of the user's mail client before accepting mail. This may be by username and password (as some webmail services provide) or through other means.

This course will only cover how to configure a basic unauthenticated null client.

The outgoing mail relay then uses DNS to look up the **MX** record identifying the mail server that accepts delivery for messages sent to the recipient's domain. The relay then uses SMTP on port 25/TCP to transmit the email to that server.

The recipient's mail service may provide a POP3 or IMAP server, such as **Dovecot** or **Cyrus**, to allow a dedicated mail client to download their messages. Frequently, the mail service provides a web-based interface, allowing clients to use a web browser as a mail client.

The following image illustrates how an email client retrieves incoming mail from an IMAP server and sends outgoing mail through an SMTP server. The mail client on **serverX.example.com** fetches incoming mails from the IMAP server **imapX.example.com**. Outgoing mails are sent to **smtpX.example.com**. An **MX** DNS record defines **smtpX.example.com** as the responsible mail server for the **desktopX.example.com** domain.

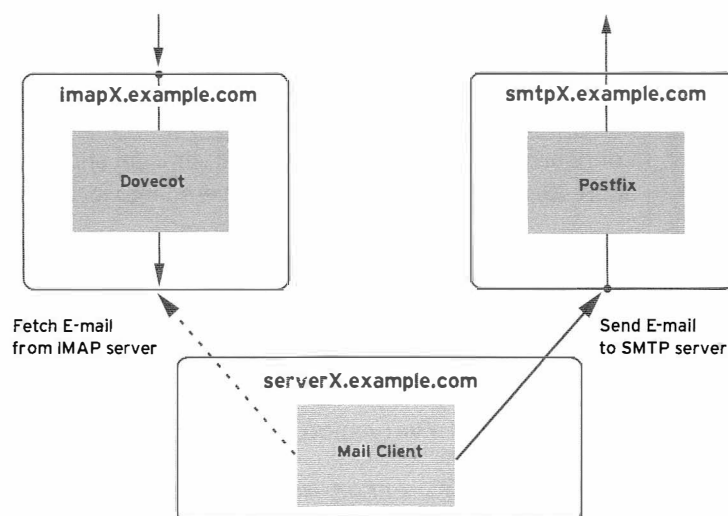


Figure 6.1: Email client communication

Postfix

Postfix is a powerful but easy-to-configure mail server. It is the default mail server in Red Hat Enterprise Linux 7. Postfix is provided by the *postfix* RPM package. It is a modular program made up of several cooperating programs. Its components are controlled by the **master** process.

The main configuration file of the **postfix** mail server is **/etc/postfix/main.cf**.



Note

There are other configuration files present in the **/etc/postfix** directory. One of the important files is **/etc/postfix/master.cf**, which controls what subservices are started.

Important Postfix Configuration Settings

Setting	Purpose
inet_interfaces	Controls which network interfaces Postfix listens on for incoming and outgoing messages. If set to loopback-only , Postfix listens only on 127.0.0.1 and ::1 . If set to all , Postfix listens on all network interfaces. One or more host names and IP addresses, separated by white space, can be listed. Default: inet_interfaces = localhost
myorigin	Rewrite locally posted email to appear to come from this domain. This helps ensure responses return to the correct domain the mail server is responsible for. Default: myorigin = \$myhostname
relayhost	Forward all messages to the mail server specified that are supposed to be sent to foreign mail addresses. Square brackets around the host name suppress the MX record lookup. Default: relayhost =
mydestination	Configure which domains the mail server is an end point for. Email addressed to these domains are delivered into local mailboxes. Default: mydestination = \$myhostname, localhost, \$mydomain, localhost
local_transport	Determine how email addressed to \$mydestination should be delivered. By default, set to local:\$myhostname , which uses the local mail delivery agent to deliver incoming mail to the local message store in /var/spool/mail . Set this to error: error message , e.g., local_transport = error: local delivery disabled , to disable local delivery completely. Default: local_transport = local:\$myhostname
mynetworks	Allow relay through this mail server from a comma-separated list of IP addresses and networks in CIDR notation to anywhere, without further authentication. If the mynetworks setting is not explicitly set in /etc/postfix/main.cf , it will be filled automatically using the setting for mynetworks_style . The default for mynetworks_style is subnet , meaning that all subnets in which the server has an IP address will be added to mynetworks . This is often not a desired situation, especially in situations where the server has an external IP address. It is recommended that a mynetworks setting gets added manually, or mynetworks_style is set to host . Default: mynetworks = 127.0.0.0/8 [::1]/128

The configuration file **/etc/postfix/main.cf** can be edited in two ways: by hand using a text editor such as **vim**, or it can be edited using the **postconf** utility. The **postconf** command

allows for querying by individual or all settings, modifying settings, querying defaults, or showing all settings that differ from the built-in defaults:

- Run the **postconf** command without any parameter to query all settings from the **/etc/postfix/main.cf** configuration file:

```
[root@serverX ~]# postconf
2bounce_notice_recipient = postmaster
access_map_defer_code = 450
access_map_reject_code = 554
address_verify_cache_cleanup_interval = 12h
address_verify_default_transport = $default_transport
...
```

- Query a particular set of options by listing them after the **postconf** command, separated by white space. Run the following to list the **inet_interfaces** and **myorigin** options with their corresponding values:

```
[root@serverX ~]# postconf inet_interfaces myorigin
inet_interfaces = loopback-only
myorigin = $myhostname
```



Note

If a value in **/etc/postfix/main.cf** starts with a dollar sign (\$), it is not a literal value, but instead points to the value of a different setting. In the previous example, the **myorigin** setting will have the same value as the **myhostname** setting. Using this syntax can simplify maintenance, since the value only has to be updated in one place.

- Run the following to add new or change existing options in the **/etc/postfix/main.cf** configuration file: **postconf -e 'setting = value'** If there was already a setting by that name in the configuration file, it will be updated to the new value; otherwise, it will be added to the bottom of the configuration file.

Run the following to change the **myorigin** setting to rewrite the domain part of the **FROM:** E-mail address to **example.com**:

```
[root@serverX ~]# postconf -e 'myorigin = example.com'
```



Important

The **postfix** service requires a **reload** or **restart** after the changes have been made to **/etc/postfix/main.cf**.



Note

When troubleshooting email, a log of all mail-related operations is kept in the **systemd** journal and **/var/log/maillog**, which includes information of any mail server-related actions.

The **postqueue -p** command displays a list of any outgoing mail messages that have been queued. To attempt to deliver all queued messages again immediately, run the **postqueue -f** command; otherwise, Postfix will attempt to resend them about once an hour until they are accepted or expire.

Postfix null client configuration

Remember, to act as a null client, Postfix and the RHEL system must be configured so that the following things are true:

- The **sendmail** command and programs that use it forward all emails to an existing outbound mail relay for delivery
- The local Postfix service does not accept local delivery for *any* email messages
- Users may run mail clients on the null client to read and send emails.

The following diagram illustrates how a null client setup works. For sending mails, the null client on **serverX** delivers all messages to the corporate SMTP mail server **smtpX.example.com**.

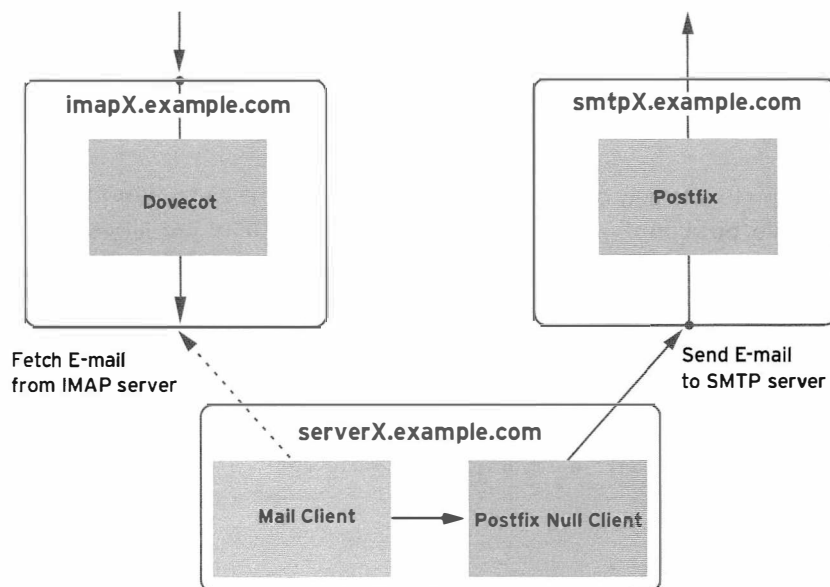


Figure 6.2: Null client communication



Note

A complete overview of all settings that are adjustable in `/etc/postfix/main.cf` and their explanations can be found in the `postconf(5)` man page. To access this man page, use the `man 5 postconf` command. Omitting the `5` will display the man page for the `postconf` command rather than the configuration guide.

Configure Postfix as null client

In this example, please follow along with these steps while your instructor demonstrates how to configure Postfix on `serverX.example.com` as a null client that uses the `smtpX.example.com` mail server as a relay host. All mails passed on by the local null client have their sender address domain rewritten to `desktopX.example.com`.

Adjust the configuration of the Postfix mail server on your serverX system to act as a null client that forwards all messages to the corporate mail server on `smtpX.example.com`, which is responsible for the `desktopX.example.com` domain.

1. Adjust the `relayhost` directive to point to the corporate mail server. The host name of the corporate mail server needs to be enclosed in square brackets to prevent an MX record lookup with the DNS server.

```
[root@serverX ~]# postconf -e "relayhost=[smtpX.example.com]"
```

2. Configure the Postfix mail server to only relay emails from the local system.
 1. Let the Postfix mail server listen only on the loopback interface for emails to deliver.

```
[root@serverX ~]# postconf -e "inet_interfaces=loopback-only"
```

2. Change the configuration of the null client so that mails originating from the `127.0.0.0/8` IPv4 network and the `:::1/128` IPv6 network are forwarded to the relay host by the local null client.

```
[root@serverX ~]# postconf -e "mynetworks=127.0.0.0/8 :::1/128"
```

3. Configure Postfix so all outgoing mails have their sender domain rewritten to the company domain `desktopX.example.com`.

```
[root@serverX ~]# postconf -e "myorigin=desktopX.example.com"
```

4. Prohibit the Postfix mail server from delivering any messages to local accounts.
 1. Configure the null client to not act as an end point for any mail domain. Mails where the recipient is a local email account are not accepted for local delivery. The `mydestination` option needs to be set to an empty value to achieve this.

```
[root@serverX ~]# postconf -e "mydestination="
```


2. Configure the local null client to not sort any mails into mailboxes on the local system. Local email delivery is turned off.

```
[root@serverX ~]# postconf -e "local_transport=error: local delivery disabled"
```

5. Restart the local **postfix** null client.

```
[root@serverX ~]# systemctl restart postfix
```

The following table shows a summary of the `/etc/postfix/main.cf` configuration file settings to configure Postfix on **serverX.example.com** as a null client that uses the **smtpX.example.com** mail server as a relay host. All messages that are passed on by the local null client have their sender address domain rewritten to **desktopX.example.com**.

Null Client Postfix Settings

Directive	Null Client (serverX.example.com)
<i>inet_interfaces</i>	<code>inet_interfaces = loopback-only</code>
<i>myorigin</i>	<code>myorigin = desktopX.example.com</code>
<i>relayhost</i>	<code>relayhost = [smtpX.example.com]</code>
<i>mydestination</i>	<code>mydestination =</code>
<i>local_transport</i>	<code>local_transport = error: local delivery disabled</code>
<i>mynetworks</i>	<code>mynetworks = 127.0.0.0/8, [::1]/128</code>



References

`postconf(1)`, `postconf(5)`, `mail(1)`, and `mutt(1)` man pages

Practice: Configuring Send-only Email Service

Guided exercise

In this lab, you will configure a local mail server as a null client.

Resources:

Machines:	desktopX and serverX
------------------	----------------------

Outcome:

Configure a local mail server acting as a null client that forwards all messages to a central server for delivery.

Before you begin...

- Reset the serverX system.
- Reset the desktopX system.
- Log into and set up your desktop system.

```
[student@desktopX ~]$ lab smtp-nullclient setup
```

- Log into and set up your server system after the desktop system setup is completed.

```
[student@serverX ~]$ lab smtp-nullclient setup
```

- Become **root** on your **serverX** system.

```
[student@serverX ~]$ sudo -i
```

Adjust the configuration of the Postfix mail server on your desktopX system to act as a null client that forwards all mail to the corporate mail server on **smtpX.example.com**. Successfully grade the desktop and the server system with **lab smtp-nullclient grade**.

- 1. Switch to user **root** on serverX and point the **relayhost** directive to the corporate mail server. The host name of the corporate mail server needs to be enclosed in square brackets to prevent an MX record lookup with the DNS server.

```
[root@serverX ~]# postconf -e "relayhost=[smtpX.example.com]"
```

- 2. Configure the Postfix mail server to only relay mail from the local system.
 - 2.1. Let the Postfix mail server only listen for emails to deliver on the loopback interface. Add the **inet_interfaces=loopback-only** directive to the **/etc/postfix/main.cf** configuration file.

```
[root@serverX ~]# postconf -e "inet_interfaces=loopback-only"
```

- 2.2. Change the Postfix configuration so that only messages that originate from the **127.0.0.0/8** IPv4 network and the **[::1]/128** IPv6 network are forwarded to the relay host by the null client on serverX.

```
[root@serverX ~]# postconf -e "mynetworks=127.0.0.0/8 [::1]/128"
```

- 3. Configure the null client so that all outgoing messages have their sender domain rewritten to the company domain **desktopX.example.com**.

```
[root@serverX ~]# postconf -e "myorigin=desktopX.example.com"
```

- 4. Prohibit the Postfix mail server from delivering any mail to local accounts.

- 4.1. Configure the null client to forward all mail to the relay server. The **mydestination** option needs to be set to an empty value to achieve this.

```
[root@serverX ~]# postconf -e "mydestination="
```

- 4.2. Prevent the local null client from sorting any mail into mailboxes on the serverX system.

```
[root@serverX ~]# postconf -e "local_transport=error: local delivery disabled"
```

- 5. Restart the local **postfix** null client on serverX.

```
[root@serverX ~]# systemctl restart postfix
```

- 6. Open a new terminal on serverX and test the null client configuration by sending an email with **Subject: serverX null client** and content **null client test** to **student@desktopX.example.com** with the **mail** command on serverX. The mail command uses **/usr/sbin/sendmail**, provided by Postfix to transfer email.

```
[student@serverX ~]$ mail -s "serverX null client" student@desktopX.example.com
null client test
.
EOT
```

- 7. Verify the mail arrived at the specified recipient by using the **mutt** command-line mail client to access the IMAP server running on **imapX.example.com** with the **IMAPS** protocol. The IMAP server is accessible for user **student** with password **student**.

- 7.1. Connect to the IMAP server running on **imapX.example.com** with the **mutt** command-line mail client with the IMAPS protocol.

```
[student@serverX ~]$ mutt -f imaps://imapX.example.com
```

- 7.2. Log into the IMAP server with user name **student** and password **student**.

```
User name at imapX.example.com: student
```

```
Password for student@imapX.example.com: student
```

- 7.3. Verify the message arrived in the mail account of **student**, then exit mutt with the **q** key.

```
q:Quit d:Del u:Undel s:Save m:Mail r:Reply g:Group ?:Help
...
  1  F May 20                (0.4K) serverX null client
q
```

8. Run the grading script on both serverX and desktopX:

```
[student@serverX ~]$ lab smtp-nullclient grade
```

```
[student@desktopX ~]$ lab smtp-nullclient grade
```

Lab: Configuring Email Transmission

Performance checklist

In this lab, you will configure a local mail server as a null client.

Resources:**Machines:**

desktopX and serverX

Outcome:

Configure a local mail server acting as a null client that forwards all messages to a central server for delivery.

Before you begin...

- Reset the serverX system.
- Reset the desktopX system.
- Log into and set up your desktop system.

```
[student@desktopX ~]$ lab smtp-nullclient setup
```

- Log into and set up your server system after the desktop system setup is completed.

```
[student@serverX ~]$ lab smtp-nullclient setup
```

- Become **root** on your **serverX** system.

```
[student@serverX ~]$ sudo -i
```

Configure the Postfix mail server on serverX.example.com as a null client which relays mail through smtpX.example.com, using desktopX.example.com as your organization's domain name on outgoing email.

Verify that the mail server is working by using an IMAPS-capable mail client to retrieve a test email from imapX.example.com (as user and mail recipient **student** with IMAP password **student**).

Successfully grade the desktop and the server system with **lab smtp-nullclient grade**.

1. Switch to user **root** on serverX and point the **relayhost** directive to the corporate mail server. The host name of the corporate mail server needs to be enclosed in square brackets to prevent the null client from performing an MX record lookup using DNS.
2. Configure the Postfix mail server to only relay mail from the local system.
3. Configure the null client so that all outgoing messages have their sender domain rewritten to the company domain **desktopX.example.com**.
4. Prohibit the Postfix mail server from delivering any mail to local accounts.

5. Restart the local **postfix** null client on serverX.
6. Open a new terminal on serverX and test the null client configuration by using the **mail** command to send an email with **Subject: serverX null client** and content **null client test** to **student@desktopX.example.com**.
7. Verify the mail arrived at the specified recipient by using the **mutt** command-line mail client to access the IMAP server running on **imapX.example.com** with the **IMAPS** protocol. The IMAP server is accessible for user **student** with password **student**.
8. Run the grading script on both serverX and desktopX:

```
[student@serverX ~]$ lab smtp-nullclient grade
```

```
[student@desktopX ~]$ lab smtp-nullclient grade
```

Solution

In this lab, you will configure a local mail server as a null client.

Resources:

Machines:	desktopX and serverX
------------------	----------------------

Outcome:

Configure a local mail server acting as a null client that forwards all messages to a central server for delivery.

Before you begin...

- Reset the serverX system.
- Reset the desktopX system.
- Log into and set up your desktop system.

```
[student@desktopX ~]$ lab smtp-nullclient setup
```

- Log into and set up your server system after the desktop system setup is completed.

```
[student@serverX ~]$ lab smtp-nullclient setup
```

- Become **root** on your **serverX** system.

```
[student@serverX ~]$ sudo -i
```

Configure the Postfix mail server on serverX.example.com as a null client which relays mail through smtpX.example.com, using desktopX.example.com as your organization's domain name on outgoing email.

Verify that the mail server is working by using an IMAPS-capable mail client to retrieve a test email from imapX.example.com (as user and mail recipient **student** with IMAP password **student**).

Successfully grade the desktop and the server system with **lab smtp-nullclient grade**.

1. Switch to user **root** on serverX and point the **relayhost** directive to the corporate mail server. The host name of the corporate mail server needs to be enclosed in square brackets to prevent the null client from performing an MX record lookup using DNS.

```
1.1. [root@serverX ~]# postconf -e "relayhost=[smtpX.example.com]"
```

2. Configure the Postfix mail server to only relay mail from the local system.
 - 2.1. Let the Postfix mail server only listen for emails to deliver on the loopback interface. Add the **inet_interfaces=loopback-only** directive to the **/etc/postfix/main.cf** configuration file.

```
[root@serverX ~]# postconf -e "inet_interfaces=loopback-only"
```

- 2.2. Change the Postfix configuration so that only messages that originate from the **127.0.0.0/8** IPv4 network and the **:::1/128** IPv6 network are forwarded to the relay host by the null client on serverX.

```
[root@serverX ~]# postconf -e "mynetworks=127.0.0.0/8 [:::1]/128"
```

3. Configure the null client so that all outgoing messages have their sender domain rewritten to the company domain **desktopX.example.com**.

- 3.1.

```
[root@serverX ~]# postconf -e "myorigin=desktopX.example.com"
```

4. Prohibit the Postfix mail server from delivering any mail to local accounts.

- 4.1. Configure the null client to forward all mail to the relay server. The **mydestination** option needs to be set to an empty value to achieve this.

```
[root@serverX ~]# postconf -e "mydestination="
```

- 4.2. Prevent the local null client from sorting any mail into mailboxes on the serverX system.

```
[root@serverX ~]# postconf -e "local_transport=error: local delivery disabled"
```

5. Restart the local **postfix** null client on serverX.

- 5.1.

```
[root@serverX ~]# systemctl restart postfix
```

6. Open a new terminal on serverX and test the null client configuration by using the **mail** command to send an email with **Subject: serverX null client** and content **null client test** to **student@desktopX.example.com**.

- 6.1.

```
[student@serverX ~]$ mail -s "serverX null client" student@desktopX.example.com
null client test
.
EOT
```

7. Verify the mail arrived at the specified recipient by using the **mutt** command-line mail client to access the IMAP server running on **imapX.example.com** with the **IMAPS** protocol. The IMAP server is accessible for user **student** with password **student**.

- 7.1. Connect to the IMAP server running on **imapX.example.com** with the **mutt** command-line mail client with the IMAPS protocol.

```
[student@serverX ~]$ mutt -f imaps://imapX.example.com
```

- 7.2. Log into the IMAP server with user name **student** and password **student**.


```
User name at imapX.example.com: student
```

```
Password for student@imapX.example.com: student
```

- 7.3. Verify the message arrived in the mail account of **student**, then exit mutt with the **q** key.

```
q:Quit d:Del u:Undel s:Save m:Mail r:Reply g:Group ?:Help
...
 1  F May 20                (0.4K) serverX null client
q
```

8. Run the grading script on both serverX and desktopX:

```
[student@serverX ~]$ lab smtp-nullclient grade
```

```
[student@desktopX ~]$ lab smtp-nullclient grade
```

Summary

Configuring Send-only Email Service

In this section, students learned how to configure the Postfix mail server to act as a null client that forwards local emails to a central SMTP server.



CHAPTER 7

PROVIDING REMOTE BLOCK STORAGE

Overview	
Goal	To provide and use networked iSCSI block devices as remote disks.
Objectives	<ul style="list-style-type: none">• Explain at a high level how iSCSI is used to provide remote access to block devices.• Provide remote access using a local disk as a LUN of an iSCSI storage target.• Access remote storage using an iSCSI initiator and prepare it for use.
Sections	<ul style="list-style-type: none">• iSCSI Concepts (and Practice)• Providing iSCSI Targets (and Practice)• Accessing iSCSI Storage (and Practice)
Lab	<ul style="list-style-type: none">• Providing Block-based Storage

iSCSI Concepts

Objectives

After completing this section, students should be able to:

- Describe the network stack of iSCSI protocol components.
- Describe the topology distinctions between Fibre Channel and iSCSI.
- Describe the distinction between block I/O and file I/O.

Introduction to iSCSI

The Internet Small Computer System Interface (iSCSI) is a TCP/IP-based protocol for emulating a SCSI high-performance local storage bus over IP networks, providing data transfer and management to remote block storage devices. As a storage area network (SAN) protocol, iSCSI extends SANs across local and wide area networks (LANs, WANs, and the Internet), providing location-independent data storage retrieval with distributed servers and arrays.

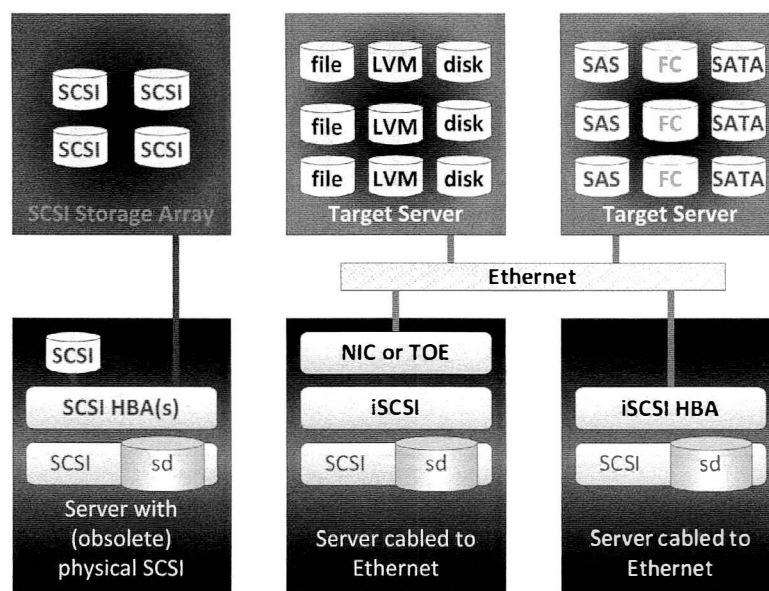


Figure 7.1: SCSI and iSCSI block storage topologies

The SCSI protocol suite provides the Command Descriptor Block (CDB) command set over a device bus communication protocol. The original SCSI topology used physical cabling with a 20-meter limitation for all devices per channel (cabled bus). Devices used unique numeric target IDs (0-7, or 0-15 with dual channel). Physical SCSI disks and cabling were obsolete by popular implementation of Fibre Channel (FC), which retained the SCSI CDB command set but replaced the disk and bus communication with protocols for longer and faster optical cabling.

The iSCSI protocol also retains the CDB command set, performing bus communication between iSCSI systems that is encapsulated over standard TCP/IP. iSCSI servers emulate SCSI devices using files, logical volumes, or disks of any type as the underlying storage (*backstore*) presented as *targets*. An iSCSI service is typically implemented in software above either an operating system TCP/IP stack or a TCP offload engine (TOE), a specialized Ethernet network interface

card (NIC) that includes the TCP/IP network layers to increase performance. iSCSI can also be hardware-implemented as a host bus adapter (HBA) for a greater performance increase.

Enterprise-grade SANs require dedicated traffic infrastructure. FC's independent optical cabling and switches guarantees isolation. iSCSI should also be implemented on cabling that is independent of standard LAN traffic, since performance can degrade due to bandwidth congestion on shared networks. Both Ethernet and FC now offer copper and optical cabling options, allowing network consolidation combined with traffic classification.

Storage area network traffic is typically unencrypted, since physical server-to-storage cabling is normally enclosed within secure data centers. For WAN security, iSCSI and Fibre Channel over Ethernet (FCoE) can utilize Internet Protocol Security (IPSec), a protocol suite for securing IP network traffic. Select networking hardware (preferred NICs, TOEs, and HBAs) can provide encryption. iSCSI offers Challenge-Handshake Authentication Protocol (CHAP) usernames and passwords as an authentication mechanism to limit connectivity between chosen initiators and targets.

Until recently, iSCSI was not considered an enterprise-grade storage option, primarily due to the use of slower 100 and 1000 Mb/s Ethernet, compared to FC's 1 and 4 Gb/s optical infrastructure. With current 10 or 40 Gb/s Ethernet and 8, 10, 16, or 20 Gb/s FC, and pending 100 Gb/s Ethernet and 32 or 128 Gb/s FC, bandwidth availability is now similar for both.




Figure 7.2: Fibre Channel block storage topologies

The use of iSCSI extends a SAN beyond the limits of local cabling, facilitating storage consolidation in local or remote data centers. Because iSCSI structures are logical, new storage allocations are made using only software configuration, without the need for additional cable or physical disks. iSCSI also simplifies data replication, migration and disaster recovery using multiple remote data centers.

iSCSI fundamentals

The iSCSI protocol functions in a familiar client-server configuration. Client systems configure *initiator* software to send SCSI commands to remote server storage *targets*. Accessed iSCSI targets appear on the client system as local, unformatted SCSI block devices, identical to devices connected with SCSI cabling, FC direct attached, or FC switched fabric.

iSCSI Component Terminology

Term	Description
initiator	An iSCSI client, typically available as software but also implemented as iSCSI HBAs. Initiators must be given unique names (see IQN).
target	An iSCSI storage resource, configured for connection from an iSCSI server. Targets must be given unique names (see IQN). A target provides one or more numbered block devices called logical units (see LUN). An iSCSI server can provide many targets concurrently.
ACL	An Access Control List (entry), an access restriction using the node IQN (commonly the iSCSI Initiator Name) to validate access permissions for an initiator.
discovery	Querying a target server to list configured targets. Target use requires an additional access steps (see login).
IQN	An iSCSI Qualified Name, a worldwide unique name used to identify both initiators and targets, in the mandated naming format: iqn.YYYY-MM.com.reversed.domain[:optional_string] iqn —Signifying that this name will use a domain as its identifier. YYYY-MM —The first month in which the domain name was owned. com.reversed.domain —The reversed domain name of the organization creating this iSCSI name. optional_string —An optional, colon-prefixed string assigned by the domain owner as desired while remaining worldwide unique. It may include colons to separate organization boundaries.
login	Authenticating to a target or LUN to begin client block device use.
LUN	A Logical Unit Number, numbered block devices attached to and available through a target. One or more LUNs may be attached to a single target, although typically a target provides only one LUN.
node	Any iSCSI initiator or iSCSI target, identified by its IQN.
portal	An IP address and port on a target or initiator used to establish connections. Some iSCSI implementations use <i>portal</i> and <i>node</i> interchangeably.
TPG	Target Portal Group, the set of interface IP addresses and TCP ports to which a specific iSCSI target will listen. Target configuration (e.g., ACLs) can be added to the TPG to coordinate settings for multiple LUNs.

iSCSI uses ACLs to perform *LUN masking*, managing the accessibility of appropriate targets and LUNs to initiators. Access to targets may also be limited with CHAP authentication. iSCSI ACLs are similar to FC's use of device worldwide numbers (WWNs) for *soft zoning* management restrictions. Although FC switch-level compulsory port restriction (*hard zoning*) has no comparable iSCSI mechanism, Ethernet VLANs could provide similar isolation security.

Unlike local block devices, iSCSI network-accessed block devices are discoverable from many remote initiators. Typical local file systems (e.g., ext4, XFS, btrfs) do not support concurrent multisystem mounting, which can result in significant file system corruption. Clustered systems resolve multiple system access by use of the Global File System (GFS2), designed to provide distributed file locking and concurrent multinode file system mounting.

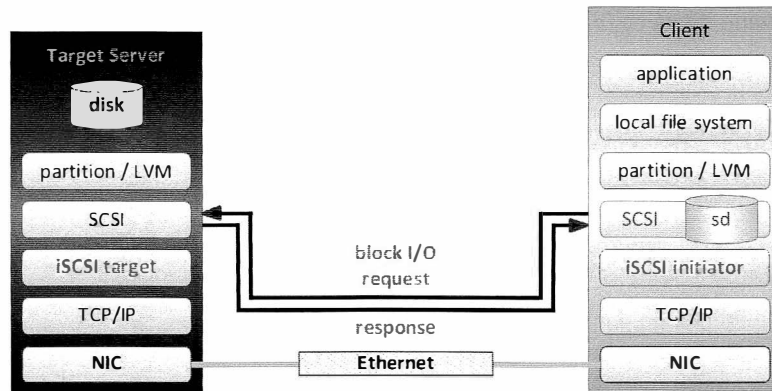


Figure 7.3: Block I/O network stack components

An attached iSCSI block device appears as a local SCSI block device (**sd**x) for use underneath a local file system, swap space, or a raw database installation as diagrammed previously. See the following for contrast with the use of network file server protocols (e.g., NFS, SMB), providing file I/O from remote file systems to local applications on multiple client systems concurrently.

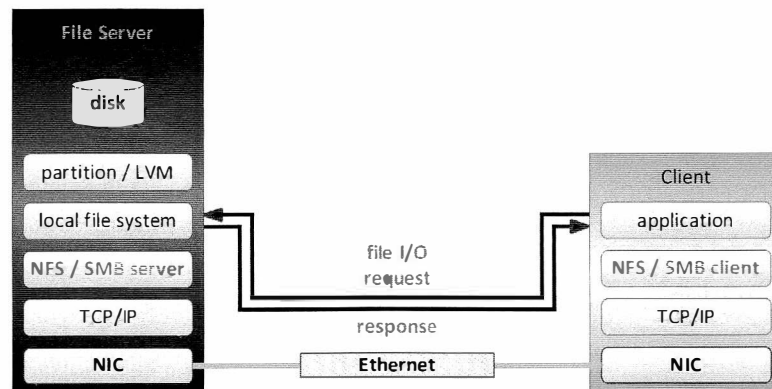


Figure 7.4: File I/O network stack components



References

Open-iSCSI project

<http://www.open-iscsi.org/>

Technical Committee T10 - SCSI Storage Interfaces

<http://www.t10.org/>

RFC 3270 - Internet Small Computer Systems Interface (iSCSI)

<http://www.ietf.org/rfc/rfc3270.txt>

Practice: iSCSI Concepts

Quiz

Match the following items to their counterparts in the table.

IQN	LUN	WWN	initiator	node	portal	target
-----	-----	-----	-----------	------	--------	--------

Description	Term
Unique name to identify individual iSCSI targets and initiators.	
Unique number to identify individual Fibre Channel port and nodes.	
Storage resource on an iSCSI server.	
Storage resource block device on an iSCSI server.	
iSCSI client implemented in either software or hardware.	
A single iSCSI initiator or target.	
A single IP connection address on an initiator or target.	

Solution

Match the following items to their counterparts in the table.

Description	Term
Unique name to identify individual iSCSI targets and initiators.	IQN
Unique number to identify individual Fibre Channel port and nodes.	WWN
Storage resource on an iSCSI server.	target
Storage resource block device on an iSCSI server.	LUN
iSCSI client implemented in either software or hardware.	initiator
A single iSCSI initiator or target.	node
A single IP connection address on an initiator or target.	portal

Providing iSCSI Targets

Objectives

After completing this section, students should be able to:

- Prepare a Red Hat Enterprise Linux server to provide the iSCSI target service.
- Build iSCSI targets and LUNs for network access.
- Limit access to iSCSI targets and LUNs.

iSCSI target overview

In original SCSI protocol terminology, a target is a single connectible storage or output device uniquely identified on a SCSI bus. In iSCSI, in which the SCSI bus is emulated across an IP network, a target can be a dedicated physical device in a network-attached storage enclosure or an iSCSI software-configured logical device on a networked storage server. A target, like HBAs and initiators, is an end point in SCSI bus communication, passing command descriptor blocks (CDB) to request or provide storage transactions.

To provide access to the storage or output device, a target is configured with one or more logical unit numbers (LUNs). In iSCSI, LUNs appear as the target's sequentially numbered disk drives, although targets typically have only one LUN. An initiator performs SCSI negotiation with a target to establish a connection to the LUN. The LUN responds as an emulated SCSI disk block device, which can be used in raw form or formatted with a client-supported file system.



Warning

Do not mount single-system file systems to more than one system at a time. iSCSI allows shared target and LUN access from multiple initiator nodes, requiring the use of cluster-capable file systems such as GFS2. Mounting standard file systems designed for local, single-system access (e.g., btrfs, ext3, ext4, FAT32, HPFS+, NTFS, XFS, ZFS) from more than one system concurrently will cause file system corruption.

iSCSI provides for *LUN masking* by using ACLs to restrict LUN accessibility to specific initiators. Except when shared access is intended, ACLs can ensure that only a designated client node can login to a specific target. On the target server, ACLs can be set at the TPG level to secure groups of LUNs, or set individually per LUN.

iSCSI target configuration

Target server configuration demonstration

targetcli is both a command-line utility and an interactive shell in which to create, delete, and configure iSCSI target components. Target stack objects are grouped into a hierarchical tree of objects, allowing easy navigation and contextual configuration. Familiar Linux commands are used in this shell: **cd**, **ls**, **pwd**, and **set**.

The **targetcli** shell also supports **TAB** completion. Administrators can use **TAB** to either complete partially typed commands or view a list of acceptable key-words at the current location in a command.

In this example, please follow along with the steps below while your instructor demonstrates using **targetcli** to configure a target server.

1. Install **targetcli** if needed.

```
[root@server0 ~]# yum -y install targetcli
```

2. Run **targetcli** with no options to enter interactive mode.

```
[root@server0 ~]# targetcli
/> ls
```

```
o- / ..... [....]
  o- backstores ..... [....]
    | o- block ..... [Storage Objects: 0]
    | o- fileio ..... [Storage Objects: 0]
    | o- pscsi ..... [Storage Objects: 0]
    | o- ramdisk ..... [Storage Objects: 0]
  o- iscsi ..... [Targets: 0]
  o- loopback ..... [Targets: 0]
```

Figure 7.5: Initial targetcli blank configuration

3. Create backing storage (backstores). There are several types of backing storage.
 - **block** - A block device defined on the server. A disk drive, disk partition, a logical volume, multipath device, any device files defined on the server that are of type **b**.
 - **fileio** - Create a file, of a specified size, in the filesystem of the server. This method is similar to using image files to be the storage for virtual machine disk images.
 - **pscsi** - Physical SCSI. Permits passthrough to a physical SCSI device connected to the server. This backstore type is not typically used.
 - **ramdisk** - Create a ramdisk device, of a specified size, in memory on the server. This type of storage will not store data persistently. When the server is rebooted, the ramdisk definition will return when the target is instantiated, but all data will have been lost.

Examples include using an existing logical volume, a disk partition, and a new file at a specified size. The backstores display as **deactivated**.

```
> cd /backstores/
/backstores> block/ create block1 /dev/iSCSI_vg/disk1_lv
Created block storage object block1 using /dev/iSCSI_vg/disk1_lv.
/backstores> block/ create block2 /dev/vdb2
Created block storage object block2 using /dev/vdb2.
/backstores> fileio/ create file1 /root/disk1_file 100M
Created fileio file1 with size 104857600
/backstores> ls
o- backstores ..... [....]
  o- block ..... [Storage Objects: 2]
    | o- block1 ..... [/dev/iSCSI_vg/disk1_lv (100.0MiB) write-thru deactivated]
    | o- block2 ..... [/dev/vdb2 (1.0GiB) write-thru deactivated]
  o- fileio ..... [Storage Objects: 1]
    | o- file1 ..... [/root/disk1_file (100.0MiB) write-back deactivated]
  o- pscsi ..... [Storage Objects: 0]
  o- ramdisk ..... [Storage Objects: 0]
```

4. Create an IQN for the target. This step will also create a default TPG underneath the IQN.

```

/backstores> cd /iscsi/
/iscsi> create iqn.2014-06.com.example:remotedisk1
Created target iqn.2014-06.com.example:remotedisk1.
Created TPG 1.
/iscsi> ls
o- iscsi ..... [Targets: 1]
  o- iqn.2014-06.com.example:remotedisk1 ..... [TPGs: 1]
    o- tpg1 ..... [no-gen-acls, no-auth]
      o- acls ..... [ACLs: 0]
      o- luns ..... [LUNs: 0]
      o- portals ..... [Portals: 0]

```

An administrator can use the **create** without specifying the IQN to create. **targetcli** will generate an IQN similar to the following: **iqn.2003-01.org.linux-iscsi.server0.x8664:sn.69b30d2cfd01**. Specifying the IQN value provides the ability for an administrator to use a meaningful namespace for their IQNs.

5. In the TPG, create an ACL for the client node to be used later. Because the global **auto_add_mapped_luns** parameter is set to true (default), any existing LUNs in the TPG are mapped to each ACL as it is created.

```

/iscsi> cd iqn.2014-06.com.example:remotedisk1/tpg1/
/iscsi/iqn.20...sk1/tpg1> acls/ create iqn.2014-06.com.example:desktop0
Created NODE ACL for iqn.2014-06.com.example:desktop0
/iscsi/iqn.20...server0/tpg1> ls
o- tpg1 ..... [no-gen-acls, no-auth]
  o- acls ..... [ACLs: 1]
    | o- iqn.2014-06.com.example:desktop0 ..... [Mapped LUNs: 0]
  o- luns ..... [LUNs: 0]
  o- portals ..... [Portals: 0]

```

This ACL configures the target to only accept initiator connections from a client presenting **iqn.2014-06.com.example:desktop0** as its initiator IQN, also known as the initiator name.

6. In this TPG, create a LUN for each existing backstores. This step also activates each backstore. Because acls exist for the TPG, they will be automatically assigned to each LUN created.

```

/iscsi/iqn.20...:server0/tgp1> luns/ create /backstores/block/block1
Created LUN 0.
Created LUN 0->0 mapping in node ACL iqn.2014-06.com.example:desktop0
/iscsi/iqn.20...:server0/tgp1> luns/ create /backstores/block/block2
Created LUN 1.
Created LUN 1->1 mapping in node ACL iqn.2014-06.com.example:desktop0
/iscsi/iqn.20...:server0/tgp1> luns/ create /backstores/fileio/file1
Created LUN 2.
Created LUN 2->2 mapping in node ACL iqn.2014-06.com.example:desktop0
/iscsi/iqn.20...:server0/tgp1> ls
o- tgp1 ..... [no-gen-acls, no-auth]
  o- acls ..... [ACLs: 1]
    | o- iqn.2014-06.com.example:desktop0 ..... [Mapped LUNs: 3]
    |   o- mapped_lun0 ..... [lun0 block/block1 (rw)]
    |   o- mapped_lun1 ..... [lun1 block/block2 (rw)]
    |   o- mapped_lun2 ..... [lun2 fileio/file1 (rw)]
  o- luns ..... [LUNs: 3]
    | o- lun0 ..... [block/block1 (/dev/iSCSI_vg/disk1_lv)]
    | o- lun1 ..... [block/block2 (/dev/vdb2)]
    | o- lun2 ..... [fileio/file1 (/root/disk1_file)]
  o- portals ..... [Portals: 0]

```

Having three LUNs assigned to a target means that when the initiator connects to the target, it will receive three new SCSI devices.

7. Still inside the TPG, create a portal configuration to designate the listening IP address and ports. Create a portal using the system's public network interface. Without specifying a TCP port to use, the portal creation will default to the standard iSCSI port (3260).

```

/iscsi/iqn.20...:server0/tgp1> portals/ create 172.25.0.11
Using default IP port 3260
Created network portal 172.25.0.11:3260
/iscsi/iqn.20...:server0/tgp1> ls
o- tgp1 ..... [no-gen-acls, no-auth]
  o- acls ..... [ACLs: 1]
    | o- iqn.2014-06.com.example:desktop0 ..... [Mapped LUNs: 3]
    |   o- mapped_lun0 ..... [lun0 block/block1 (rw)]
    |   o- mapped_lun1 ..... [lun1 block/block2 (rw)]
    |   o- mapped_lun2 ..... [lun2 fileio/file1 (rw)]
  o- luns ..... [LUNs: 3]
    | o- lun0 ..... [block/block1 (/dev/iSCSI_vg/disk1_lv)]
    | o- lun1 ..... [block/block2 (/dev/vdb2)]
    | o- lun2 ..... [fileio/file1 (/root/disk1_file)]
  o- portals ..... [Portals: 1]
    o- 172.25.0.11:3260 ..... [OK]

```

If the IP is not specified with the portal creation, an IP of 0.0.0.0 will be used. This will permit connections on all network interfaces defined on the server.

8. View the entire configuration, then exit **targetcli**. **targetcli** will automatically save upon exit. The resulting persistent configuration file is stored in JavaScript Object Notation (JSON) format.

```

/iscsi/iqn.20...:server0/tgp1> cd /
/> ls
o- / ..... [....]
  o- backstores ..... [....]
    | o- block ..... [Storage Objects: 2]
    | | o- block1 ..... [/dev/iSCSI_vg/disk1_lv (100.0MiB) write-thru activated]
    | | o- block2 ..... [/dev/vdb2 (1.0GiB) write-thru activated]
    | o- fileio ..... [Storage Objects: 1]
    | | o- file1 ..... [/root/disk1_file (100.0MiB) write-back activated]
    | o- pscsi ..... [Storage Objects: 0]
    | o- ramdisk ..... [Storage Objects: 0]
  o- iscsi ..... [Targets: 1]
    | o- iqn.2014-06.com.example:remotedisk1 ..... [TPGs: 1]
    |   o- tgp1 ..... [no-gen-acls, no-auth]
    |     o- acls ..... [ACLs: 1]
    |       | o- iqn.2014-06.com.example:desktop0 ..... [Mapped LUNs: 3]
    |         | o- mapped_lun0 ..... [lun0 block/block1 (rw)]
    |         | o- mapped_lun1 ..... [lun1 block/block2 (rw)]
    |         | o- mapped_lun2 ..... [lun2 fileio/file1 (rw)]
    |     o- luns ..... [LUNs: 3]
    |       | o- lun0 ..... [block/block1 (/dev/iSCSI_vg/disk1_lv)]
    |       | o- lun1 ..... [block/block2 (/dev/vdb2)]
    |       | o- lun2 ..... [fileio/file1 (/root/disk1_file)]
    |     o- portals ..... [Portals: 1]
    |       o- 172.25.0.11:3260 ..... [OK]
  o- loopback ..... [Targets: 0]
/> exit
Global pref auto_save_on_exit=true
Last 10 configs saved in /etc/target/backup.
Configuration saved to /etc/target/saveconfig.json

```

9. Add a port exemption to the default firewall for port 3260, the standard iSCSI port.

```

[root@server0 ~]# firewall-cmd --add-port=3260/tcp
[root@server0 ~]# firewall-cmd --add-port=3260/tcp --permanent

```

10. Enable the **target.service** systemd unit. The **target.service** will recreate the target configuration from the json file at boot. If this step is skipped, any configured targets will work until the machine is rebooted; however, after a reboot, no targets will be offered by the server.

```

[root@server0 ~]# systemctl enable target

```

Authentication

In addition to ACL node verification, password-based authentication can be implemented. Authentication can be required during the iSCSI discovery phase. Authentication can be unidirectional or bidirectional.

CHAP authentication does not use strong encryption for the passing of credentials. While CHAP does offer an additional factor of authentication besides having a correctly configured initiator name, configured in an ACL, it should not be considered secure. If security of iSCSI data is a concern, controlling the network side of the protocol is a better method to assure security. Providing a dedicated, isolated network, or vlans to pass the iSCSI traffic will be a more secure implementation of the protocol.

Command-line mode

In the demonstration, **targetcli** was run in interactive mode, but **targetcli** can also be used to execute a series of commands via the command-line. In the following example, **targetcli** will be used to create a backstore device, an IQN, and activate a portal. The example, as written, will not present a usable target, but is meant to demonstrate several actions executed with **targetcli**. This method could be used to script target configuration. At the end of the list of commands, a **saveconfig** command is executed. Unlike the interactive use of **targetcli**, command-line mode will not save the configuration as a json file until the **saveconfig** command is used.

```
[root@server0 ~]# targetcli /backstores/block create block1 /dev/vdb
Created block storage object block1 using /dev/vdb.
[root@server0 ~]# targetcli /iscsi create iqn.2014-06.com.example:remotedisk1
Created target iqn.2014-06.com.example:remotedisk1.
Created TPG 1.
[root@server0 ~]# targetcli /iscsi/iqn.2014-06.com.example:remotedisk1/tpg1/portals
create 172.25.0.11
Using default IP port 3260
Created network portal 172.25.0.11:3260.
[root@server0 ~]# targetcli saveconfig
Last 10 configs saved in /etc/target/backup.
Configuration saved to /etc/target/saveconfig.json
```

R

References

targetcli(8) man page

LINUX I/O Targetcli
<http://linux-iscsi.org/wiki/Targetcli>

Practice: Providing iSCSI Targets

Guided exercise

In this lab, you will configure a RHEL server to become an iSCSI target server, including opening firewall access, creating a backing store, and setting target and LUN access parameters.

Resources:	
Files:	/etc/target/saveconfig.json
Machines:	serverX

Outcomes:

The creation and configuration of an iSCSI target block device with ACL-validated access.

Before you begin...

- Reset your **serverX** system.
- Log in to your **serverX** system and become **root**.

```
[student@serverX ~]$ sudo -i
```

- 1. Prepare the server system to become an iSCSI target server by installing the **target** configuration utility, starting the **target** service and opening the firewall for the iSCSI server port.

- 1.1. Install the *targetcli* RPM.

```
[root@serverX ~]# yum -y install targetcli
```

- 1.2. Enable and start the **target** service.

```
[root@serverX ~]# systemctl enable target; systemctl start target
```

- 1.3. Open the iSCSI server port on the firewall as a permanent change. Reload the configuration for immediate use.

```
[root@serverX ~]# firewall-cmd --permanent --add-port=3260/tcp
[root@serverX ~]# firewall-cmd --reload
```

- 2. Create the physical disk structure to use as a backing store device for a target to be created in a later step. Partition a disk, build a volume group, and create a logical volume.

- 2.1. Create a 1 GB partition using the second physical disk on the server system.

```
[root@serverX ~]# fdisk /dev/vdb
```

Use **fdisk** to create a new primary partition with a size of 1 GB. Use the partition tag appropriate for use as a Linux LVM partition ("**8e**"). Remember to write before exiting. This exercise uses the expected partition name **/dev/vdb1**.

- 2.2. Create a logical volume manager volume group named **iSCSI_vg** using the partition created in the previous step.

```
[root@serverX services]# pvcreate /dev/vdb1
[root@serverX services]# vgcreate iSCSI_vg /dev/vdb1
[root@serverX services]# vgdisplay iSCSI_vg
```

- 2.3. Create a 100 MiB logical volume named **disk1_lv** in the new volume group.

```
[root@serverX ~]# lvcreate -n disk1_lv -L 100m iSCSI_vg
[root@serverX ~]# lvdisplay iSCSI_vg/disk1_lv
```

- 3. Go into **targetcli**'s interactive mode to configure the iSCSI target.

```
[root@serverX ~]# targetcli
```

- 4. Configure the existing **/dev/iSCSI_vg/disk1_lv** logical volume as a block-type backing store. Use the name **serverX.disk1** for the backstore, replacing **X** with your system number.

```
/> /backstores/block/ create serverX.disk1 /dev/iSCSI_vg/disk1_lv
```

- 5. Create a unique iSCSI Qualified Name (IQN) for the target. The target will be **iqn.2014-06.com.example:serverX**.

```
/> /iscsi create iqn.2014-06.com.example:serverX
```

Creating the IQN name automatically created a default target portal group named **tpg1**.

- 6. Create an ACL for the client node (initiator). The initiator will be connecting with its initiator name set to **iqn.2014-06.com.example:desktopX**.

```
/> /iscsi/iqn.2014-06.com.example:serverX/tpg1/acls/ create
iqn.2014-06.com.example:desktopX
```

- 7. Create a LUN under the target, the LUN should use the previously defined backing storage device named **serverX.disk1**.

```
/> /iscsi/iqn.2014-06.com.example:serverX/tpg1/luns create /backstores/block/
serverX.disk1
```

- 8. Configure a portal for the target to listen on 172.25.X.11, port 3260.

```
/> /iscsi/iqn.2014-06.com.example:serverX/tpg1/portals create 172.25.X.11
```

- 9. View, verify, and save the target server configuration.
 - 9.1. View and verify the configuration.

```

/> ls
o- / ..... [....]
  o- backstores ..... [....]
    | o- block ..... [Storage Objects: 1]
    | | o- serverX.disk1 .....
    | | .....[/dev/iSCSI_vg/disk1_lv (100.0MiB) write-thru activated]
    | o- fileio ..... [Storage Objects: 0]
    | o- pscsi ..... [Storage Objects: 0]
    | o- ramdisk ..... [Storage Objects: 0]
  o- iscsi ..... [Targets: 1]
    | o- iqn.2014-06.com.example:serverX ..... [TPGs: 1]
    | | o- tpg1 ..... [no-gen-acls, no-auth]
    | | | o- acls ..... [ACLs: 1]
    | | | | o- iqn.2014-06.com.example:desktopX ..... [Mapped LUNs: 1]
    | | | | | o- mapped_lun0 ..... [lun1 block/serverX.disk1 (rw)]
    | | | | o- luns ..... [LUNs: 1]
    | | | | | o- lun0 ..... [block/serverX.disk1 (/dev/iSCSI_vg/disk1_lv)]
    | | | o- portals ..... [Portals: 1]
    | | | | o- 172.25.X.11:3260 ..... [OK]
  o- loopback ..... [Targets: 0]

```

- 9.2. Exit and save the configuration to the default `/etc/target/saveconfig.json`.

```

/> exit

```



Important

The outcome of the practice will be used in the next practice. Therefore, do not reset your serverX host after the completion of this practice.



Note

A practical method to test this configuration is to install the iSCSI client tools package and perform a discovery from a separate client system. This is presented in the *Accessing iSCSI Storage* exercise.

Although an iSCSI client configuration can be installed on this target server, an administrator must ensure that any iSCSI node records created during test discoveries are deleted prior to target discovery from client systems. These procedures are presented in the *Accessing iSCSI Storage* section.

Accessing iSCSI Storage

Objectives

After completing this section, students should be able to:

- Prepare a Red Hat Enterprise Linux client to access an iSCSI target service.
- Access and log into a remote iSCSI portal, target, and LUN.
- Practice iSCSI persistence through the use and deletion of persistent node records.

iSCSI initiator introduction

In Red Hat Enterprise Linux an iSCSI initiator is typically implemented in software, and functions similar to a hardware iSCSI HBA to access targets from a remote storage server. Using a software-based iSCSI initiator requires connecting to an existing Ethernet network of sufficient bandwidth to carry the expected storage traffic.

iSCSI can also be implemented using a hardware initiator that includes the required protocols in a dedicated host bus adapter. iSCSI HBAs and TCP offload engines (TOE), which include the TCP network stack on an Ethernet NIC, move the processing of iSCSI or TCP overhead and Ethernet interrupts to hardware, easing the load on system CPUs.

Configuring an iSCSI client initiator requires installing the *iscsi-initiator-utils* package, which includes the **iscsi** and **iscsid** services and the **/etc/iscsi/iscsid.conf** and **/etc/iscsi/initiatorname.iscsi** configuration files.

As an iSCSI node, the client requires a unique IQN. The default **/etc/iscsi/initiatorname.iscsi** file contains a generated IQN using Red Hat's domain. Administrators typically reset the IQN to their own domain and an appropriate client system string.

The **/etc/iscsi/iscsid.conf** file contains default settings for node records created during new target discovery. Settings include iSCSI timeouts, retry parameters, and authentication usernames and passwords. Changing this file requires restarting the **iscsi** service.

```
[root@desktopX ~]# systemctl restart iscsi
```

To be able to discover targets, install the *iscsi-initiator-utils* package, then enable and start the **iscsi** service. Targets must be discovered before device connection and use. The discovery process stores target node information and settings in the **/var/lib/iscsi/nodes** directory, using defaults from **/etc/iscsi/iscsid.conf**. Since the same target can exist on multiple portals, node records are stored for each portal. Perform discovery with the following command:

```
[root@desktopX ~]# iscsiadm -m discovery -t sendtargets -p target_server[:port]
172.25.X.11:3260,1 iqn.2014-06.com.example:serverX
```

In discovery mode, the **sendtargets** request returns only targets with access configured for this initiator. The port number can be omitted when the target server is configured on default port 3260. Upon discovery, a node record is written to **/var/lib/iscsi/nodes** and used for subsequent logins. To use the listed target, log in using the following command:

- 2.3. Identify the newly available block device created by the iSCSI target login.

```
[root@desktopX ~]# lsblk
[root@desktopX ~]# tail /var/log/messages
```

- 3. Browse the connection information about the target portal, connection, and parameters used by the connected device. Locate the node record.

- 3.1. List the targets recognized by the **iscsi** service.

```
[root@desktopX ~]# iscsiadm -m session -P 3
iSCSI Transport Class version 2.0-870
version 6.2.0-873-21
Target: iqn.2014-06.com.example:serverX
  Current Portal: 172.25.X.11:3260,1
  Persistent Portal: 172.25.X.11:3260,1
  *****
  Interface:
  *****Resources
  Iface Name: default
  Iface Transport: tcp
  Iface Initiatorname: iqn.2014-06.com.example:desktopX
  Iface IPaddress: 172.25.X.10
  Iface HWaddress: <empty>
  Iface Netdev: <empty>
  SID: 3
  iSCSI Connection State: LOGGED IN
  iSCSI Session State: LOGGED_IN
  Internal iscsid Session State: NO CHANGE
  *****
  Negotiated iSCSI params:
  *****
  HeaderDigest: None
  DataDigest: None
  MaxRecvDataSegmentLength: 262144
  MaxXmitDataSegmentLength: 8192
  FirstBurstLength: 65536
  MaxBurstLength: 262144
  ImmediateData: Yes
  InitialR2T: Yes
  MaxOutstandingR2T: 1
  *****
  Attached SCSI devices:
  *****
  Host Number: 5 State: running
  scsi5 Channel 00 Id 0 Lun: 0
  scsi5 Channel 00 Id 0 Lun: 1
  Attached scsi disk sdb State: running
```

- 3.2. Change directory to the location of the iSCSI node records for the remainder of this exercise. Locate the persistent node record for the new iSCSI target.

```
[root@desktopX ~]# cd /var/lib/iscsi/nodes
[root@desktopX nodes]$ ls -lR
```

- 3.3. View the connection parameters defaults in an iSCSI node record.

```
[root@desktopX nodes]# less
iqn.2014-06.com.example:serverX/172.25.X.11,3260,1/default
```

- 4. Connect to and disconnect from the target, using both manual commands and methods that utilize the node record(s) to connect automatically.
 - 4.1. Disconnect the iSCSI block device by logging out of the iSCSI target. Confirm that the corresponding iSCSI block device has disappeared.



Important

In this exercise, the block device was not formatted or otherwise used. If the device is in use, unmount it properly before continuing.

```
[root@desktopX nodes]# iscsiadm -m node -T iqn.2014-06.com.example:serverX
-p 172.25.X.11 -u
[root@desktopX nodes]$ lsblk
```

- 4.2. Confirm that the node record for the disconnected iSCSI target still exists.

```
[root@desktopX nodes]# ls -lR
```

- 4.3. Restart the **iscsi** client service. Confirm that the iSCSI block device returns. The **iscsi** performs logins for all node records found.

```
[root@desktopX nodes]# systemctl restart iscsi
[root@desktopX nodes]# lsblk
```

- 4.4. Disconnect the iSCSI block device again by logging out of the iSCSI target. Additionally, delete the node record using the proper command. Confirm that the iSCSI block device has again disappeared.

```
[root@desktopX nodes]# iscsiadm -m node -T iqn.2014-06.com.example:serverX
-p 172.25.X.11 -u
[root@desktopX nodes]# iscsiadm -m node -T iqn.2014-06.com.example:serverX
-p 172.25.X.11 -o delete
[root@desktopX nodes]# lsblk
```

- 4.5. Confirm that the node record for the disconnected iSCSI target no longer exists.

```
[root@desktopX nodes]# ls -lR
```

- 4.6. Restart the **iscsi** client service. Confirm that the iSCSI block device does not return since there is no node record to trigger the target login.

```
[root@desktopX nodes]# systemctl restart iscsi
[root@desktopX nodes]# lsblk
```

- 5. Rediscover the target and confirm that the discovery step alone creates the node record(s).
 - 5.1. Rediscover the configured iSCSI targets provided by a specific iSCSI target server portal, but do *not* log into the target.

```
[root@desktopX nodes]# iscsiadm -m discovery -t st -p 172.25.X.11  
172.25.X.11:3260,1 iqn.2014-06.com.example:serverX
```

- 5.2. Confirm that the node record for the iSCSI target is created by discovery, even before login has occurred.

```
[root@desktopX nodes]# ls -lR
```

- 5.3. Clean up by deleting the node record again. Confirm that the node record no longer exists. When finished, return to your home directory.

```
[root@desktopX nodes]# iscsiadm -m node -T iqn.2014-06.com.example:serverX  
-p 172.25.X.11 -o delete  
[root@desktopX nodes]# ls -lR  
[root@desktopX nodes]# cd ~
```

Lab: Providing Block-based Storage

Performance checklist

In this lab, you will build a target server and connect to a target from a client node, then use the resulting block devices to create a file system.

Resources:	
Files:	<ul style="list-style-type: none"> • /etc/target/saveconfig.json • /etc/iscsi/initiatorname.iscsi
Machines:	serverX and desktopX

Outcomes:

An iSCSI target exported from your **serverX** machine, limited with ACLs so that it can only be used by your **desktopX**, used by your **desktopX** system with a persistent mount.

Before you begin...

- Reset your **serverX** system.
- Reset your **desktopX** system.

You have been asked to create a new 1 GiB iSCSI target on your **serverX** machine. This target should be called **iqn.2014-06.com.example:serverX**. This target should only be available to clients with an initiatorname of **iqn.2014-06.com.example:desktopX**.

Your **desktopX** machine should persistently log in to this target, and persistently mount a (new) XFS file system created on that target on **/iscsidisk**.

1. On your **serverX** machine create a new **1 GiB** partition on your secondary hard disk for use as a backing store.
2. Install all needed packages on **serverX**, then start and enable all relevant services.
3. Create a new iSCSI target called **iqn.2014-06.com.example:serverX**, using your new partition as a backing store, and limited only to clients with the initiatorname **iqn.2014-06.com.example:desktopX**.
4. Open all relevant ports for iSCSI on the firewall on your **serverX** machine.
5. On your **desktopX** machine verify that the *iscsi-initiator-utils* package is installed, then set the initiatorname to **iqn.2014-06.com.example:desktopX**.
6. Discover and login to the new target. Verify the device node for the new target.
7. Create an XFS file system on the new disk, then persistently mount it on **/iscsidisk**.

Solution

In this lab, you will build a target server and connect to a target from a client node, then use the resulting block devices to create a file system.

Resources:	
Files:	<ul style="list-style-type: none"> • /etc/target/saveconfig.json • /etc/iscsi/initiatorname.iscsi
Machines:	serverX and desktopX

Outcomes:

An iSCSI target exported from your **serverX** machine, limited with ACLs so that it can only be used by your **desktopX**, used by your **desktopX** system with a persistent mount.

Before you begin...

- Reset your **serverX** system.
- Reset your **desktopX** system.

You have been asked to create a new 1 GiB iSCSI target on your **serverX** machine. This target should be called **iqn.2014-06.com.example:serverX**. This target should only be available to clients with an initiatorname of **iqn.2014-06.com.example:desktopX**.

Your **desktopX** machine should persistently log in to this target, and persistently mount a (new) XFS file system created on that target on **/iscsidisk**.

1. On your **serverX** machine create a new **1 GiB** partition on your secondary hard disk for use as a backing store.

```
1.1. [student@serverX ~]$ sudo fdisk /dev/vdb
Welcome to fdisk (util-linux 2.23.2).

Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Device does not contain a recognized partition table
Building a new DOS disklabel with disk identifier 0xaa23f397.

Command (m for help): n
Partition type:
   p   primary (0 primary, 0 extended, 4 free)
   e   extended
Select (default p): p
Partition number (1-4, default 1): Enter
First sector (2048-20971519, default 2048): Enter
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519): +1G
Partition 1 of type Linux and of size 1 GiB is set

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

2. Install all needed packages on **serverX**, then start and enable all relevant services.

- 2.1. Install the *targetcli* package.

```
[student@serverX ~]$ sudo yum install targetcli
```

- 2.2. Enable and start the **target.service** service.

```
[student@serverX ~]$ sudo systemctl enable target.service
[student@serverX ~]$ sudo systemctl start target.service
```

3. Create a new iSCSI target called **iqn.2014-06.com.example:serverX**, using your new partition as a backing store, and limited only to clients with the initiatorname **iqn.2014-06.com.example:desktopX**.

- 3.1.

```
[student@serverX ~]$ sudo targetcli
targetcli shell version 2.1.fb34
Copyright 2011-2013 by Datera, Inc and others.
For help on commands, type 'help'.

/> /backstores/block create disk1 /dev/vdb1
Created block storage object disk1 using /dev/vdb1.
/> /iscsi create iqn.2014-06.com.example:serverX
Created target iqn.2014-06.com.example:serverX.
Created TPG 1.
/> /iscsi/iqn.2014-06.com.example:serverX/tpg1/acls create
    iqn.2014-06.com.example:desktopX
Created Node ACL for iqn.2014-06.com.example:desktopX
/> /iscsi/iqn.2014-06.com.example:serverX/tpg1/luns create /backstores/block/
disk1
Created LUN 0.
Created LUN 0->0 mapping in node ACL iqn.2014-06.com.example:desktopX
/> /iscsi/iqn.2014-06.com.example:serverX/tpg1/portals create 172.25.X.11 3260
Using default IP port 3260
Created network portal 172.25.X.11:3260.
/> ls
o- / ..... [....]
  o- backstores ..... [....]
    | o- block ..... [Storage Objects: 1]
    | | o- disk1 ..... [/dev/vdb1 (1.0GiB) write-thru activated]
    | o- fileio ..... [Storage Objects: 0]
    | o- pscsi ..... [Storage Objects: 0]
    | o- ramdisk ..... [Storage Objects: 0]
  o- iscsi ..... [Targets: 1]
    | o- iqn.2014-06.com.example:serverX ..... [TPGs: 1]
    | | o- tpg1 ..... [no-gen-acls, no-auth]
    | | | o- acls ..... [ACLs: 1]
    | | | | o- iqn.2014-06.com.example:desktopX ..... [Mapped LUNs: 1]
    | | | | | o- mapped_lun0 ..... [lun0 block/disk1 (rw)]
    | | | o- luns ..... [LUNs: 1]
    | | | | o- lun0 ..... [block/disk1 (/dev/vdb1)]
    | | | o- portals ..... [Portals: 1]
    | | | | o- 172.25.X.11:3260 ..... [OK]
  o- loopback ..... [Targets: 0]
/> saveconfig
Last 10 configs saved in /etc/target/backup.
Configuration saved to /etc/target/saveconfig.json
/> exit
```

4. Open all relevant ports for iSCSI on the firewall on your **serverX** machine.

4.1.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --add-port=3260/tcp
[student@serverX ~]$ sudo firewall-cmd --reload
```

5. On your **desktopX** machine verify that the *iscsi-initiator-utils* package is installed, then set the initiatorname to **iqn.2014-06.com.example:desktopX**.

- 5.1. Verify that the *iscsi-initiator-utils* package is installed.

```
[student@serverX ~]$ sudo yum install iscsi-initiator-utils
```

- 5.2. Set your initiatorname to **iqn.2014-06.com.example:desktopX**. Edit **/etc/iscsi/initiatorname.iscsi** so that it reads:

```
InitiatorName=iqn.2014-06.com.example:desktopX
```



Note

After changing the initiatorname when a target has already been discovered do not forget to restart the **iscsid.service** service.

6. Discover and login to the new target. Verify the device node for the new target.

- 6.1. Discover the target.

```
[student@serverX ~]$ sudo iscsiadm -m discovery -t st -p serverX.example.com
```

- 6.2. Log in to the target.

```
[student@serverX ~]$ sudo iscsiadm -m node -T iqn.2014-06.com.example:serverX -l
```

- 6.3. Verify the device node for the iSCSI disk.

```
[student@serverX ~]$ sudo iscsiadm -m session -P3 | grep Attached
Attached SCSI devices:
Attached scsi disk sda      State: running
```

7. Create an XFS file system on the new disk, then persistently mount it on **/iscsidisk**.

- 7.1. Create the file system, make sure to use the device node you discovered in the previous step, then find the UUID for the new files system.

```
[student@serverX ~]$ sudo mkfs -t xfs /dev/sda
[student@serverX ~]$ sudo blkid /dev/sda
```

- 7.2. Create the mountpoint.

```
[student@serverX ~]$ sudo mkdir /iscsidisk
```

- 7.3. Add the following line to **/etc/fstab**, where **XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXX** is the UUID you discovered for your new filesystem:

```
UUID=XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX /iscsidisk xfs _netdev 0 2
```

- 7.4. Mount the file system and verify.

```
[student@serverX ~]$ sudo mount /iscsidisk
[student@serverX ~]$ df -h /iscsidisk
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda        1014M  33M  982M   4% /iscsidisk
```

Summary

iSCSI Concepts

In this section, students learned how to:

- Describe the history of the SCSI protocol and adaptation for use over networks.
- Describe block storage cabling topologies.
- Define iSCSI components and terms.

Providing iSCSI Targets

In this section, students learned how to:

- Identify the steps to install and configure an iSCSI target server.
- Describe the components for building a single iSCSI target.
- Describe the components for providing access security for each iSCSI target.

Accessing iSCSI Storage

In this section, students learned how to:

- Describe the installation and configuration of an iSCSI initiator.
- Outline the steps to access and login to an iSCSI LUN.
- Describe how iSCSI persistent connections are implemented.



CHAPTER 8

PROVIDING FILE-BASED STORAGE

Overview	
Goal	To provide NFS exports and SMB file shares to specific systems and users.
Objectives	<ul style="list-style-type: none">• Export file systems to client systems using NFS, controlling access by IP address.• Export file systems to clients using NFS, controlling access with Kerberos and using labeled NFS.• Share file systems with clients using SMB, controlling access by username and password.• Mount an SMB share with the <code>multiuser</code> mount option, using password-based authentication and <code>cifscreds</code> to control access
Sections	<ul style="list-style-type: none">• Exporting NFS File Systems (and Practice)• Protecting NFS Exports (and Practice)• Providing SMB File Shares (and Practice)• Performing a Multiuser SMB Mount (and Practice)
Lab	<ul style="list-style-type: none">• Providing File-based Storage

Exporting NFS File Systems

Objectives

After completing this section, students should be able to export file systems to client systems using NFS, controlling access by IP address.

What is NFS?

The Network File System (NFS) is a network file system commonly used by UNIX systems and network-attached storage devices to allow multiple clients to share access to files over the network. It provides access to shared directories or files from client systems.



Warning

The NFS protocol transmits data in clear text over the network. Furthermore, the server relies on the client to identify users. It is not recommended to export directories with sensitive information without the use of Kerberos authentication and encryption, which is covered later in this section.

NFS exports

A NFS server installation requires the *nfs-utils* package to be installed. It provides all necessary utilities to export a directory with NFS to clients. The configuration file for the NFS server exports is the `/etc/exports` file.

The `/etc/exports` file lists the directory to share to client hosts over the network and indicates which hosts or networks have access to the export.



Note

Instead of adding the information required for exporting directories to the `/etc/exports` file, a newly created file named `*.exports` can be added to the `/etc/exports.d/` directory holding the configuration of exports.



Warning

Exporting the same directory with NFS and Samba is not supported on Red Hat Enterprise Linux 7, because NFS and Samba use different file locking mechanisms, which can cause file corruption.

One or more clients can be listed, separated by a space, as a:

- DNS-resolvable host name, like `server0.example.com` in the following example, where the `/myshare` directory is exported and can be mounted by `server0.example.com`.

```
/myshare server0.example.com
```

- DNS-resolvable host name with the wildcards ***** for multiple characters and/or **?** for a single character. The following example allows all subdomains in the **example.com** domain to access the NFS export.

```
/myshare *.example.com
```

- DNS-resolvable host name with character class lists in square brackets. In this example, the hosts **server0.example.com**, **server1.example.com**, ..., and **server20.example.com** have access to the NFS export.

```
/myshare server[0-20].example.com
```

- IPv4 address. The following example allows access to the **/myshare** NFS share from the **172.25.11.10** IP address.

```
/myshare 172.25.11.10
```

- IPv4 network. This example shows an **/etc/exports** entry, which allows access to the NFS-exported directory **/myshare** from the **172.25.0.0/16** network.

```
/myshare 172.25.0.0/16
```

- IPv6 address without square brackets. The following example allows the client with IPv6 address **2000:472:18:b51:c32:a21** access to the NFS-exported directory **/myshare**.

```
/myshare 2000:472:18:b51:c32:a21
```

- IPv6 network without square brackets. This example allows the IPv6 network **2000:472:18:b51::/64** access to the NFS export.

```
/myshare 2000:472:18:b51::/64
```

- A directory can be exported to multiple hosts simultaneously by specifying multiple targets with their options, separated by spaces, following the directory to export.

```
/myshare *.example.com 172.25.0.0/16
```

Optionally, there can be one or more export options specified in round brackets as a comma-separated list, directly followed by each client definition. Commonly used export options are:

- **ro**, read-only: the default setting when nothing is specified. It is allowed to explicitly specify it with an entry. Restricts the NFS clients to read files on the NFS share. Any write operation is prohibited. The following example explicitly states the **ro** flag for the **server0.example.com** host.

```
/myshare desktop0.example.com(ro)
```


- **rw**, read-write: allows read and write access for the NFS clients. In the following example, the **desktop0.example.com** is able to access the NFS export read-only, while **server[0-20].example.com** has read-write access to the NFS share.

```
/myshare desktop0.example.com(ro) server[0-20].example.com(rw)
```

- **no_root_squash**: By default, **root** on an NFS client is treated as user **nfsnobody** by the NFS server. That is, if **root** attempts to access a file on a mounted export, the server will treat it as an access by user **nfsnobody** instead. This is a security measure that can be problematic in scenarios where the NFS export is used as **/** by a diskless client and **root** needs to be treated as **root**. To disable this protection, the server needs to add **no_root_squash** to the list of options set for the export in **/etc/exports**.

The following example allows the client **diskless.example.com** read-write and real root user access to the exported NFS directory **/myshare**.

```
/myshare diskless.example.com(rw,no_root_squash)
```



Warning

Note that this particular configuration is not secure, and would be better done in conjunction with Kerberos authentication and integrity checking, which is covered later in this section.

Configuring an NFS export

In this example, please follow along with these steps while your instructor demonstrates how to share a directory IP-based with NFS. The directory **/myshare** is on serverX and will be mounted on the desktopX system.

1. Start the NFS service on serverX with the **systemctl** command.

```
[root@serverX ~]# systemctl start nfs-server
```

2. Enable the NFS service to start at boot on serverX.

```
[root@serverX ~]# systemctl enable nfs-server
```

3. Create the directory **/myshare** to share it with NFS on the serverX system.

```
[root@serverX ~]# mkdir /myshare
```

4. Export the **/myshare** directory on serverX to the desktopX client as a read- and write-enabled share. To do that, add the following line to the **/etc/exports** file on serverX:

```
/myshare desktopX(rw)
```

5. After the changed `/etc/exports` file has been saved, apply the changes by executing `exportfs -r`.

```
[root@serverX ~]# exportfs -r
```

6. The NFS port 2049/TCP for `nfsd` must be open on the server. To configure `firewalld` to enable access to the NFS exports immediately, run:

```
[root@serverX ~]# firewall-cmd --permanent --add-service=nfs
```

7. Reload the `firewalld` rules so the new rule gets applied.

```
[root@serverX ~]# firewall-cmd --reload
```

8. Use the newly created mount point `/mnt/nfsexport` on the desktopX system to mount the NFS-exported directory.

```
[root@desktopX ~]# mkdir /mnt/nfsexport
```

9. On the desktopX system, the share can now be mounted on the newly created mount point `/mnt/nfsexport` with the `mount` command.

```
[root@desktopX ~]# mount serverX:/myshare /mnt/nfsexport
```

R

References

`nfs(5)`, `mount(8)`, `mount.nfs(8)`, `exportfs(8)`, and `exports(5)` man pages

Practice: Exporting NFS File Systems

Guided exercise

In this lab, you will use NFS to provide shared storage.

Resources:**Machines:**

desktopX and serverX

Outcomes:

The NFS server exports an IP-based NFS share on serverX. The NFS export is mounted on desktopX.

Before you begin...

- Reset the serverX system.
- Reset the desktopX system.
- Become **root** on your **serverX** system.

```
[student@serverX ~]$ sudo -i
```

- Become **root** on your **desktopX** system.

```
[student@desktopX ~]$ sudo -i
```

Configure an IP-based NFS share on serverX according to the following requirements:

- The NFS server provides newly created shared directory **/nfsshare**.
- The **/nfsshare** NFS export provides read and write access for **nfsnobody**.

Mount the share on the **/mnt/nfsshare** mount point on desktopX permanently.

Successfully grade the setup on desktopX and serverX with **lab nfsshare grade**.

1. Configure an IP-based NFS share on serverX that provides a newly created shared directory **/nfsshare** for the desktopX machine with read and write access for **nfsnobody**.

- 1.1. Start the NFS service on serverX.

```
[root@serverX ~]# systemctl start nfs-server
```

- 1.2. Enable the NFS service to start at boot on serverX.

```
[root@serverX ~]# systemctl enable nfs-server
```

- 1.3. Create the directory **/nfsshare** to be shared by NFS on the serverX system.

```
[root@serverX ~]# mkdir /nfsshare
```

- 1.4. Change the ownership on the **/nfsshare** to user **nfsnobody**, so the directory is writable by **nfsnobody**.

```
[root@serverX ~]# chown nfsnobody /nfsshare
```

- 1.5. Change the **/etc/exports** configuration file on serverX to share the newly created **/nfsshare** directory on the desktopX system with read and write access.

```
[root@serverX ~]# echo '/nfsshare desktopX(rw)' >>/etc/exports
```

- 1.6. Use the **exportfs -r** command to reload the **/etc/exports** configuration file on serverX.

```
[root@serverX ~]# exportfs -r
```

- 1.7. Configure **firewalld** to allow access to the NFS service on serverX.

```
[root@serverX ~]# firewall-cmd --permanent --add-service=nfs
```

- 1.8. Reload the **firewalld** configuration to allow access to the NFS service instantly on serverX.

```
[root@serverX ~]# firewall-cmd --reload
```

- 2. Mount the NFS export from the serverX system on the **/mnt/nfsshare** mount point on desktopX permanently.

- 2.1. Create the mount point **/mnt/nfsshare** on the desktopX system.

```
[root@desktopX ~]# mkdir /mnt/nfsshare
```

- 2.2. Create the required entry in **/etc/fstab** to mount the exported NFS share on the newly created **/mnt/nfsshare** directory on the desktopX system permanently.

```
serverX:/nfsshare /mnt/nfsshare nfs defaults 0 0
```

- 2.3. Mount the exported NFS share on the newly created **/mnt/nfsshare** directory on the desktopX system and verify the **/etc/fstab** entry works as expected.

```
[root@desktopX ~]# mount -a
```

- 2.4. Verify that the NFS share mounted at `/mnt/nfsshare` is writable on the `desktopX` system.

```
[root@desktopX ~]# touch /mnt/nfsshare/test.txt
[root@desktopX ~]# ls -l /mnt/nfsshare
total 0
-rw-r--r--. 1 nfsnobody nfsnobody 0 May  8 04:14 test.txt
```

- 3. Run the grading script on both `serverX` and `desktopX`:

```
[student@serverX ~]$ lab nfsshare grade
```

```
[student@desktopX ~]$ lab nfsshare grade
```

Protecting NFS Exports

Objectives

After completing this section, students should be able to export file systems to clients using NFS, controlling access with Kerberos and using labeled NFS.

Kerberos-enabled NFS exports

The NFS server does not require authentication and only enforces access restriction based on the IP address or host name of the client by default. To remedy this, the NFS server provides options to secure access to files using a number of methods: **none**, **sys**, **krb5**, **krb5i**, and **krb5p**. The NFS server can choose to offer a single method or multiple methods for each exported share. NFS clients must connect to the exported share using one of the methods mandated for that share, specified as a mount option **sec=method**.

Security Methods

- **none**: Anonymous access to the files, writes to the server will be allocated UID and GID of **nfsnobody**. This requires the SELinux Boolean **nfsd_anon_write** to be active.
- **sys**: File access based on standard Linux file permissions for UID and GID values. If not specified, this is the default. The NFS server trusts any UID sent by the client.
- **krb5**: Clients must prove identity using Kerberos and then standard Linux file permissions apply. UID/GID is determined based upon the Kerberos principal from the accessing user.
- **krb5i**: Adds a cryptographically strong guarantee that the data in each request has not been tampered with. UID/GID is determined based upon the Kerberos principal from the accessing user.
- **krb5p**: Adds encryption to all requests between the client and the server, preventing data exposure on the network. This will have a performance impact, but provides the most security. UID/GID is determined based upon the Kerberos principal from the accessing user.

For using any of the security options that use a Kerberos server, the **nfs-secure-server** needs to be running in addition to the **nfs-server** service on the system exporting the NFS shares. The client requires the **nfs-secure** service to run to help with negotiating Kerberos authentication.



Important

Kerberos options will require, at a minimum, a **/etc/krb5.keytab** and additional authentication configuration that is not covered in this section, joining the Kerberos realm. The **/etc/krb5.keytab** will normally be provided by the authentication or security administrator. Request a **keytab** that includes either a *host principal*, *nfs principal*, or both.

Configuring a Kerberos-enabled NFS server

In this example, please follow along with these steps while your instructor demonstrates how to export a directory with NFS using **krb5p** security. The required keytab files are provided. The NFS

server exports the **/securedexport** directory to all subdomains of **example.com**. The client mounts the NFS export on the **/mnt/securedexport** mount point.

1. Install the keytab provided at <http://classroom.example.com/pub/keytabs/serverX.keytab> on the serverX system, which will act as the NFS server.

```
[root@serverX ~]# wget -O /etc/krb5.keytab http://classroom.example.com/pub/keytabs/serverX.keytab
```

2. For NFS with Kerberos security, the **nfs-secure-server** needs to be running. Start the **nfs-secure-server** service on the serverX system.

```
[root@serverX ~]# systemctl start nfs-secure-server
```

3. Enable the **nfs-secure-server** to start at system boot on serverX.

```
[root@serverX ~]# systemctl enable nfs-secure-server
```

4. Create the directory **/securedexport** on serverX. This directory will be used as the NFS export.

```
[root@serverX ~]# mkdir /securedexport
```

5. Add the directory **/securedexport** to the **/etc/exports** file to export it with NFS. Enable **krb5p** security to secure access to the NFS share. Allow read and write access to the exported directory from all subdomains of the **example.com** domain.

```
[root@serverX ~]# echo '/securedexport *.example.com(sec=krb5p,rw)' >>/etc/exports
```

6. After the changed **/etc/exports** file has been saved, apply the changes by executing **exportfs -r**.

```
[root@serverX ~]# exportfs -r
```

7. The NFS port 2049/TCP for **nfsd** must be open on the server. To configure **firewalld** to enable access to the NFS exports immediately, run:

```
[root@serverX ~]# firewall-cmd --permanent --add-service=nfs
```

8. Reload the **firewalld** rules so the new rule gets applied.

```
[root@serverX ~]# firewall-cmd --reload
```

9. Install the provided keytab on the desktopX system, which will act as the NFS client. Mount the **krb5p**-secured share on the desktopX system:

```
[root@desktopX ~]# wget -O /etc/krb5.keytab http://classroom.example.com/pub/
keytabs/desktopX.keytab
```

10. NFS uses the **nfs-secure** service on the client side to help negotiate and manage communication with the server when connecting to Kerberos-secured shares. It must be running to use the secured NFS shares; **start** and **enable** it to ensure it is always available on the desktopX system.

```
[root@desktopX ~]# systemctl enable nfs-secure
ln -s '/usr/lib/systemd/system/nfs-secure.service' ...
[root@desktopX ~]# systemctl start nfs-secure
```



Note

The **nfs-secure** service is part of the **nfs-utils** package, which should be installed by default. If it is not installed, run:

```
[root@desktopX ~]# yum -y install nfs-utils
```

11. The mount point must exist to mount the krb5p-secured export from serverX on the desktopX system. **/mnt/secureexport** is created on the desktopX system.

```
[root@desktopX ~]# mkdir /mnt/secureexport
```

12. The exported directory now can be mounted on the desktopX system with krb5p security enabled.

```
[root@desktopX ~]# mount -o sec=krb5p serverX:/secureexport /mnt/secureexport
```

13. To clean up for the next demonstration, the **/mnt/secureexport** mount point needs to be unmounted on the desktopX system.

```
[root@desktopX ~]# umount /mnt/secureexport
```

SELinux and labeled NFS

SELinux offers additional security by locking down the capabilities of services provided in Red Hat Enterprise Linux. By default, NFS mounts have the SELinux context **nfs_t**, independent of the SELinux context they have on the server that provides the export.

This behavior can be changed on the client side by using the mount option **context="selinux_context"**. The following example mounts the NFS export and enforces the SELinux context: **system_u:object_r:public_content_rw_t:s0**:

```
[root@serverX ~]# mount -o context="system_u:object_r:public_content_rw_t:s0" serverX:/
myshare /mnt/nfsexport
```


The NFS server can be forced to properly export the SELinux context of a share by switching to NFS version 4.2. This specification currently only exists as an Internet draft. It is already implemented in the NFS server shipped by Red Hat Enterprise Linux 7, but needs to be turned on explicitly.

To enable NFS version 4.2 on the serverX system to export the SELinux labels, change the **RPCNFSDARGS=""** line in the **/etc/sysconfig/nfs** file to:

```
RPCNFSDARGS="-v 4.2"
```

The **nfs-server** or **nfs-secure-server** respectively require a restart.

```
[root@serverX ~]# systemctl restart nfs-server
[root@serverX ~]# systemctl restart nfs-secure-server
```

On the client side, **mount -o v4.2** must be specified as the mount option.

```
[root@desktopX ~]# mount -o sec=krb5p,v4.2 serverX:/securedexport /mnt/securedexport
```

For testing purposes, a new file with the name **selinux.txt** is created in the exported directory **/securedexport**. After creation, the SELinux type is changed to **public_content_t**.

```
[root@serverX ~]# touch /securedexport/selinux.txt
[root@serverX ~]# chcon -t public_content_t /securedexport/selinux.txt
```

All SELinux labels are now properly handled by serverX and forwarded to the client system desktopX.

```
[root@desktopX ~]# ls -Z /mnt/securedexport/
-rw-r--r--. root root unconfined_u:object_r:public_content_t:s0 selinux.txt
```



Note

In a default installation of Red Hat Enterprise Linux 7, the **nfs_export_all_ro** and **nfs_export_all_rw** SELinux Booleans are both enabled. This allows the NFS daemon to read and write almost any file. To lock down the capabilities of the NFS server, disable these Booleans. For content to be readable by NFS, it should have the **public_content_t** or **nfs_t** SELinux context. For content to be both readable and writable, it should have the **public_content_rw_t** or **nfs_t** context. If the **public_content_rw_t** context is used, the **nfsd_anon_write** Boolean must be enabled to allow writes. Additional NFS-related SELinux information can be found in the **nfsd_selinux(8)** man page, which is provided by the *selinux-policy-devel* RPM package.



References

nfs(5), **mount(8)**, **mount.nfs(8)**, **exportfs(8)**, **exports(5)**, and **nfsd_selinux(8)**
man pages

Practice: Protecting NFS Exports

Guided exercise

In this lab, you will use NFS to provide shared storage protected with Kerberos and SELinux.

Resources:**Machines:**

desktopX and serverX

Outcomes:

The NFS server exports a Kerberized NFS share on serverX with SELinux labels. The NFS export is mounted on desktopX with krb5p security and SELinux labels available.

Before you begin...

- Reset the serverX system.
- Reset the desktopX system.
- Log into and set up your server system.

```
[student@serverX ~]$ lab nfskrb5 setup
```

- Log into and set up your desktop system.

```
[student@desktopX ~]$ lab nfskrb5 setup
```

- Become **root** on your **serverX** system.

```
[student@serverX ~]$ sudo -i
```

- Become **root** on your **desktopX** system.

```
[student@desktopX ~]$ sudo -i
```

Configure the NFS server on serverX to meet the following requirements:

- Share the newly created **/securenfs** directory on serverX with krb5p security.
- Allow read and write access on the share from the desktopX system.
- SELinux labels are exported.
- The NFS share is mounted on the **/mnt/secureshare** desktopX with krb5p security and exported SELinux labels.
- Preconfigured krb5 keytabs for the serverX and desktopX systems are available at:

<http://classroom.example.com/pub/keytabs/serverX.keytab>.

<http://classroom.example.com/pub/keytabs/desktopX.keytab>.

- Successfully grade the setup on desktopX and serverX with **lab nfskrb5 grade**.
- 1. Configure NFS to share the newly created **/securenfs** directory on serverX with krb5p security. Allow read and write access from the desktopX system. The SELinux labels on the shared directory are exported.
 - 1.1. Install the keytab provided at `http://classroom.example.com/pub/keytabs/serverX.keytab` on the serverX system.

```
[root@serverX ~]# wget -O /etc/krb5.keytab http://classroom.example.com/pub/keytabs/serverX.keytab
```
 - 1.2. Enable NFS version 4.2 on the serverX system to export the SELinux labels. To do that, change the **RPCNFSDARGS=""** line in the **/etc/sysconfig/nfs** file to:

```
RPCNFSDARGS="-v 4.2"
```
 - 1.3. Start the **nfs-secure-server** service on the serverX system.

```
[root@serverX ~]# systemctl start nfs-secure-server
```
 - 1.4. Enable **nfs-secure-server** to start at system boot on serverX.

```
[root@serverX ~]# systemctl enable nfs-secure-server
```
 - 1.5. Create the directory **/securenfs** on serverX.

```
[root@serverX ~]# mkdir /securenfs
```
 - 1.6. Add the directory **/securenfs** to the **/etc/exports** file to export it with NFS. Enable krb5p security to secure access to the NFS share. Allow read and write access to the exported directory from the desktopX system.

```
[root@serverX ~]# echo '/securenfs desktopX(sec=krb5p,rw)' >>/etc/exports
```
 - 1.7. Reload the **/etc/exports** file on the serverX system.

```
[root@serverX ~]# exportfs -r
```
 - 1.8. Configure **firewalld** to allow access to the NFS service on serverX.

```
[root@serverX ~]# firewall-cmd --permanent --add-service=nfs
```
 - 1.9. Reload the **firewalld** configuration to allow access to the NFS service instantly on serverX.

```
[root@serverX ~]# firewall-cmd --reload
```

- 2. Mount the krb5p-secured NFS share permanently on the **/mnt/secureshare** mount point so that all exported SELinux labels are present on the desktopX system.

- 2.1. Install the keytab provided at <http://classroom.example.com/pub/keytabs/desktopX.keytab> on the desktopX system.

```
[root@desktopX ~]# wget -O /etc/krb5.keytab http://classroom.example.com/pub/keytabs/desktopX.keytab
```

- 2.2. Start the nfs-secure service on desktopX to help with negotiating authentication with a Kerberized NFS share.

```
[root@desktopX ~]# systemctl start nfs-secure
```

- 2.3. Enable the nfs-secure service to start at system boot on desktopX.

```
[root@desktopX ~]# systemctl enable nfs-secure
```

- 2.4. Create the mount point **/mnt/secureshare** on the desktopX system.

```
[root@desktopX ~]# mkdir /mnt/secureshare
```

- 2.5. Create the entry in the **/etc/fstab** file to mount the **/securenfs** share exported by the serverX system on the **/mnt/secureshare** mount point on desktopX so that the SELinux labels from the share are shown on the mount point.

```
serverX:/securenfs /mnt/secureshare nfs defaults,v4.2,sec=krb5p 0 0
```

- 2.6. Mount the exported NFS share on the newly created **/mnt/secureshare** directory on the desktopX system and verify the **/etc/fstab** entry works as expected.

```
[root@desktopX ~]# mount -a
```

- 3. Test the setup with the newly created file **/securenfs/testfile.txt** with the content "Hello World" on the serverX machine. Set the SELinux context to **public_content_t** on the file **/securenfs/testfile.txt** on serverX. Change the ownership of the **/securenfs/testfile.txt** file to **ldapuserX:ldapuserX** and the permissions to **644**. Verify that the SELinux context and the permissions are present on the mounted share on the desktopX system. Verify that user **ldapuserX** has read and write access on the **/mnt/secureshare/testfile.txt** file on desktopX.

- ❑ 3.1. Create a new file `/securenfs/testfile.txt` with the content "Hello World" on the serverX machine.

```
[root@serverX ~]# echo "Hello World" > /securenfs/testfile.txt
```

- ❑ 3.2. Set the SELinux context to `public_content_t` on the file `/securenfs/testfile.txt` on serverX.

```
[root@serverX ~]# chcon -t public_content_t /securenfs/testfile.txt
```

- ❑ 3.3. Change the ownership of the `/securenfs/testfile.txt` file to `ldapuserX:ldapuserX` on serverX.

```
[root@serverX ~]# chown ldapuserX:ldapuserX /securenfs/testfile.txt
```

- ❑ 3.4. Change the permissions of the `/securenfs/testfile.txt` file to 644 on the serverX system.

```
[root@serverX ~]# chmod 644 /securenfs/testfile.txt
```

- ❑ 3.5. Verify the SELinux context is exported on the desktopX system and available on the mounted share.

```
[root@desktopX ~]# ls -Z /mnt/secureshare
-rw-r--r--. ldapuserX ldapuserX unconfined_u:object_r:public_content_t:s0
testfile.txt
```

- ❑ 3.6. Log into the desktopX system as `ldapuserX` with password `kerberos` by using `ssh`.

```
[root@desktopX ~]# ssh ldapuserX@desktopX
...
ldapuserX@desktopX's password: kerberos
Creating home directory for ldapuserX
```

- ❑ 3.7. Verify the file `/mnt/secureshare/testfile.txt` is writable by the Kerberos-authenticated `ldapuserX`.

```
[ldapuserX@desktopX ~]$ echo "I can write" >>/mnt/secureshare/testfile.txt
[ldapuserX@desktopX ~]$ cat /mnt/secureshare/testfile.txt
Hello World
I can write
```

- ❑ 4. Run the grading script on both serverX and desktopX:

```
[student@serverX ~]$ lab nfskrb5 grade
```

```
[student@desktopX ~]$ lab nfskrb5 grade
```

Providing SMB File Shares

Objectives

After completing this section, students should be able to share directories with clients using SMB, controlling access by username and password.

What is SMB?

Server Message Block (SMB) is the standard file-sharing protocol for Microsoft Windows servers and clients. SMB file servers can be configured in a number of different ways. One of the simplest is to configure the file servers and their clients as members of a common Windows *workgroup*, which announces servers and clients to the local subnet. The file servers each manage their own local user accounts and passwords independently. More sophisticated configurations may be members of a *Windows domain* and coordinate user authentication through a *domain controller*.

Using a software package named **Samba**, Red Hat Enterprise Linux is able to act as a server for SMB file shares. Mounting SMB file shares as a client is handled by a kernel driver and utilities included in the *cifs-utils* package.

SMB file sharing with Samba

The Samba service can share Linux file systems as SMB network file shares. This section will cover the basic configuration steps needed for a Samba server to provide a file share to the members of a Windows workgroup, managing its own users locally. It will not discuss the more complex configuration required to make the Samba server a member of a Windows domain.

The basic steps that must be performed in order to configure Samba to provide an SMB file share as a workgroup member are:

1. Install the *samba* package.
2. Prepare the permissions on the directory to be shared.
3. Configure `/etc/samba/smb.conf`.
4. Set up appropriate Linux users with NTLMv2 passwords.
5. Start Samba and open the local firewall.
6. Verify that the share can be mounted from a client.

Installing Samba

To deploy the Samba service on a Red Hat Enterprise Linux system, the *samba* package must be installed. This can be done directly, or as part of the *file-server* package group:

```
[root@serverX ~]# yum install samba
```


Preparing directories for sharing



Warning

Do not use Samba to share a directory that is also an NFS export or a mounted NFS file system. This can result in file corruption, stale file locks, or other file access issues with the share.

The directory to be shared must be created if it does not already exist.

```
[root@serverX ~]# mkdir /sharedpath
```

Users and regular permissions

The permissions which should be set on the directory will depend on who needs access to it and how it will be mounted by clients.

A client normally mounts a share by authenticating access to the SMB server as a particular user. All files on the share need to be readable (and possibly writable) by the user that is used to mount the share.

SELinux contexts and Booleans

In order for Samba to work correctly when SELinux is in enforcing mode, the directory will need to have correct SELinux contexts and certain SELinux Booleans may need to be set.

If the shared directory will only be accessed through Samba, then the directory and all its subdirectories and files should be labeled **samba_share_t**, which gives Samba read and write access. It is best to configure the SELinux policy so that **restorecon** will set this type on the share and its contents if the file system is relabeled.

For example, to configure **restorecon** so that the files in **/sharedpath** have the type **samba_share_t**, and then to relabel the directory, run:

```
[root@serverX ~]# semanage fcontext -a -t samba_share_t '/sharedpath(/.*)?'
[root@serverX ~]# restorecon -vvFR /sharedpath
restorecon reset /sharedpath context unconfined_u:object_r:default_t:s0-
>system_u:object_r:samba_share_t:s0
```



Note

Samba can also serve files labeled with the SELinux types **public_content_t** (read-only) and **public_content_rw_t** (read-write). To allow read-write access to files and directories labeled **public_content_rw_t**, the SELinux Boolean **smbd_anon_write** must also be enabled.

Configuring `/etc/samba/smb.conf`

The main configuration file for Samba is `/etc/samba/smb.conf`. This file is divided into multiple sections. Each section starts with a section name in square brackets, followed by a list of parameters set to particular values.

`/etc/samba/smb.conf` starts with a `[global]` section used for general server configuration. Subsequent sections each define a file share or printer share provided by the Samba server. Two special sections may exist, `[homes]` and `[printers]`, which have special uses. Any line beginning with either a semicolon (;) or a hash (#) character is commented out.

The `[global]` section

The `[global]` section defines the basic configuration of the Samba server. There are three things which should be configured here:

1. **workgroup** is used to specify the Windows workgroup for the server. Most Windows systems default to **WORKGROUP**, although Windows XP Home defaulted to **MSHOME**. This is used to help systems browse for the server using the NetBIOS for TCP/IP name service.

To set the workgroup to **WORKGROUP**, change the existing workgroup entry in the `/etc/samba/smb.conf` to:

```
workgroup = WORKGROUP
```

2. **security** controls how clients are authenticated by Samba. For **security = user**, clients log in with a valid username and password managed by the local Samba server. This setting is the default in `/etc/samba/smb.conf`.
3. **hosts allow** is a comma-, space-, or tab-delimited list of hosts that are permitted to access the Samba service. If it is not specified, all hosts can access Samba. If it is not specified in the `[global]` section, it can be set on each share separately. If it is specified in the `[global]` section, then it will apply to all shares, regardless of whether each share has a different setting.

Hosts can be specified by host name or by source IP address. Host names are checked by reverse-resolving the IP address of the incoming connection attempt. The full syntax of this directive is described by the `hosts_access(5)` man page.

Allowed hosts can be specified in a number of ways:

- IPv4 network/prefix: **172.25.0.0/24**
- IPv4 network/netmask: **172.25.0.0/255.255.255.0**
- If the IPv4 subnet prefix is on a byte boundary: **172.25.0.**
- IPv6 network/prefix: **[2001:db8:0:1::/64]**
- Host name: **desktop.example.com**
- All hosts ending in example.com: **.example.com**

For example, to restrict access to only the hosts from the **172.25.0.0/16** network using the *trailing dot* notation, the **hosts allow** entry in the `/etc/samba/smb.conf` configuration file would read:

```
hosts allow = 172.25.
```

To additionally allow access from all host names ending with ".example.com", the `/etc/samba/smb.conf` configuration file entry would be:

```
hosts allow = 172.25. .example.com
```

File share sections

To create a file share, at the end of `/etc/samba/smb.conf`, place the share name in brackets to start a new section for the share. Some key directives should be set in this section:

1. **path** must be set to indicate which directory to share; for example, **path = /sharedpath**.
2. **writable = yes** should be set if all authenticated users should have read-write access to the share. The default setting is **writable = no**.

If **writable = no** is set, a comma-separated **write list** of users with read-write access to the share can be provided. Users not on the list will have read-only access. Members of local groups can also be specified: **write list = @management** will permit all authenticated users who are members of the Linux group "management" to have write access.

3. **valid users**, if set, specifies a list of users allowed to access the share. Users not on the list are not allowed to access the share. However, if the list is blank, *all* users can access the share.

For example, to allow only user **fred** and members of group **management** read-only access to the share **myshare**, the section would read:

```
[myshare]
path = /sharedpath
writable = no
valid users = fred, @management
```

The [homes] section

The **[homes]** section defines a special file share, which is enabled by default. This share makes local home directories available via SMB. The share name can be specified as **homes**, in which case the Samba server will convert it to the home directory path of the authenticating user, or as a specific **username**.



Important

The **samba_enable_home_dirs** SELinux Boolean allows local Linux home directories to be shared by Samba to other systems. This needs to be enabled for **[homes]** to work (**setsebool -P samba_enable_home_dirs=on**).

The **use_samba_home_dirs** Boolean, on the other hand, allows remote SMB file shares to be mounted and used as local Linux home directories. It is easy to confuse the two options.

Validating /etc/samba/smb.conf

To verify that there are no errors in the edited **smb.conf** file, the command **testparm** is available. Run **testparm** with no arguments to verify that there are no obvious syntax errors.

```
[root@serverX ~]# testparm
Load smb config files from /etc/samba/smb.conf
rlimit_max: increasing rlimit_max (1024) to minimum Windows limit (16384)
Processing section "[random]"
Processing section "[homes]"
Processing section "[printers]"
Loaded services file OK.
Server role: ROLE_STANDALONE
Press enter to see a dump of your service definitions

[global]
server string = Samba Server Version %v
log file = /var/log/samba/log.%m
max log size = 50
idmap config * : backend = tdb
cups options = raw

[random]
comment = Test File Share
path = /srv/random

[homes]
comment = Home Directories
read only = No
browseable = No

[printers]
comment = All Printers
path = /var/spool/samba
printable = Yes
print ok = Yes
browseable = No
```



Important

The directive **read only = no** is the same as **writable = yes**, which can be confusing.

Preparing Samba users

The **security = user** setting requires a Linux account with a Samba account that has a valid NTLM password. To create a Samba-only system user, keep the Linux password locked, and set the login shell to **/sbin/nologin**. This prevents the login of the user directly, or with **ssh** on the system.

For example, to create the locked Linux account for a user **fred**:

```
[root@serverX ~]# useradd -s /sbin/nologin fred
```

The *samba-client* contains the **smbpasswd** command. It can create Samba accounts and set passwords.

```
[root@serverX ~]# yum -y install samba-client
```

If **smbpasswd** is passed a username without any options, it will attempt to *change* the account password. The root user can use it with the **-a** option to add the Samba account and set the

NTLM password. The **-x** option can be used by root to delete a Samba account and password for a user.

For example, to create a Samba account for user **fred** and assign an NTLM password:

```
[root@serverX ~]# smbpasswd -a fred
New SMB password: redhat
Retype new SMB password: redhat
...
Added user fred.
```

A more powerful tool than **smbpasswd** is also available for the **root** user, **pdbedit**. For example, **pdbedit -L** will list all users with Samba accounts configured on the system. For more information, see the **pdbedit(8)** man page.

Starting Samba

Use **systemctl** to start the Samba services immediately and enable them to start at boot time:

```
[root@serverX ~]# systemctl start smb nmb
[root@serverX ~]# systemctl enable smb nmb
```

The two services these units start, **smbd** and **nmbd**, must communicate through the local firewall. Samba's **smbd** daemon normally uses TCP/445 for SMB connections. It also listens on TCP/139 for NetBIOS over TCP backward compatibility. The **nmbd** daemon uses UDP/137 and UDP/138 to provide NetBIOS over TCP/IP network browsing support.

To configure **firewalld** to allow clients to talk to the local Samba services, run:

```
[root@serverX ~]# firewall-cmd --permanent --add-service=samba
success
[root@serverX ~]# firewall-cmd --reload
success
```



Important

Samba checks periodically to determine if **/etc/samba/smb.conf** has been changed. If the configuration file has changed, Samba automatically reloads it. This will not affect any connections already established to the Samba service, until the connection is closed or Samba is completely restarted.

The command **systemctl reload smb nmb** can be used to reload the configuration file immediately, or **systemctl restart smb nmb** to restart Samba entirely.

Mounting SMB file systems

Regular SMB mounts

The **cifs-utils** package can be used to mount SMB file shares on the local system, whether from a Samba server or a native Microsoft Windows server. By default, SMB mounts use a single set of user credentials (the *mount credentials*) for mounting the share and determining access rights to files on the share. All users on the Linux system using the mount use the same credentials to determine file access.

The **mount** command is used to mount the share. By default, the protocol used to authenticate users is NTLMv2 password hashing encapsulated in Raw NTLMSSP messages (**sec=ntlmssp**), as expected by recent versions of Microsoft Windows. The mount credentials can be provided in two ways. If mounting interactively at a shell prompt, the **username=** option can be used to specify which SMB user to authenticate as; the user will be prompted for the password. If mounting automatically, a credentials file readable only by root containing the username and password can be provided with the **credentials=** option.

For example, to mount the share **myshare** from the SMB file server **serverX**, authenticating as SMB user **fred**, who has the NTLM password **redhat**:

```
[root@desktopX ~]# mkdir /mnt/myshare
[root@desktopX ~]# mount -o username=fred //serverX/myshare /mnt/myshare
Password for fred@//serverX/myshare: redhat
```

R

References

samba(7), **smb.conf(5)**, **testparm(1)**, **mount(8)**, **mount.cifs(8)**, **smbpasswd(8)**, **pdbedit(8)**, and **samba_selinux(8)** man pages

Practice: Providing SMB File Shares

Guided exercise

In this lab, you will use SMB to provide shared storage.

Resources:

Machines:	desktopX and serverX
------------------	----------------------

Outcomes:

Share a directory with SMB on serverX according to the given requirements, then mount it on desktopX.

Before you begin...

- Reset the desktopX system.
- Reset the serverX system.
- Become **root** on your **serverX** system.

```
[student@serverX ~]$ sudo -i
```

- Become **root** on your **desktopX** system.

```
[student@desktopX ~]$ sudo -i
```

Configure a SMB share on the serverX system according to the following requirements:

- Share the newly created directory **/smbshare** with SMB.
- Members of the auxiliary group **marketing** have read and write permissions on the share.
- All users that are not member of the **marketing** group have read permission.
- The Samba server is in the **mycompany** workgroup and the share name in Samba is **smbshare**.
- Create the Samba-only user **brian**, who is part of the marketing team, with the password **redhat**.
- Create the new Samba-only user **rob** with the password **redhat**, who is not part of the marketing team.
- Successfully grade the setup on the serverX system with **lab smbshare grade**.

- 1. Deploy the required RPM packages to run the SMB service on serverX.

```
[root@serverX ~]# yum -y install samba
```

- 2. Create the auxiliary system group **marketing** and the **/smbshare** directory on serverX. The **marketing** system group owns the **/smbshare** directory. Adjust the permissions on the **/smbshare** directory to have the SGID bit set, and write is prohibited by others.

The SELinux context type on the **/smbshare** directory and all newly created files and subdirectories is **samba_share_t**.

- 2.1. Create the auxiliary system group **marketing** on the serverX system.

```
[root@serverX ~]# groupadd -r marketing
```

- 2.2. Create the **/smbshare** directory on serverX.

```
[root@serverX ~]# mkdir -p /smbshare
```

- 2.3. Change the group ownership of the **/smbshare** directory to **marketing** on the serverX system.

```
[root@serverX ~]# chgrp marketing /smbshare
```

- 2.4. Adjust the permissions on the **/smbshare** directory to have the SGID bit set, and write is prohibited by others.

```
[root@serverX ~]# chmod 2775 /smbshare
```

- 2.5. Add the directory **/smbshare**, and all files shown as follows, to the SELinux policy as a directory sharing files with SMB by setting its label to **samba_share_t**.

```
[root@serverX ~]# semanage fcontext -a -t samba_share_t '/smbshare(/.*)?'
```

- 2.6. Apply the SELinux rule for the **/smbshare** directory that was added in the previous step on serverX.

```
[root@serverX ~]# restorecon -vvFR /smbshare
restorecon reset /smbshare context unconfined_u:object_r:default_t:s0-
>system_u:object_r:samba_share_t:s0
```

- 3. Change the **/etc/samba/smb.conf** configuration file on serverX to reflect the configuration requested.

Modify or confirm the following:

```
[global]
...
workgroup = mycompany
...
security = user
passdb backend = tdbsam
```

Add a section at the end of the file as follows.


```
[smbshare]
path = /smbshare
write list = @marketing
```

- 4. Create the Samba-only user **brian**, who is part of the marketing team. The user **brian** has read and write access to the **smbshare** SMB share. A new Samba user **rob** is created, who is not part of the marketing team. The user **rob** has read access to the **smbshare** SMB share. Both newly added users have the SMB password **redhat**.

- 4.1. Install the *samba-client* RPM package because it contains **smbpasswd**.

```
[root@serverX ~]# yum -y install samba-client
```

- 4.2. Create the system user **brian** as a member of the auxiliary group **marketing** on serverX.

```
[root@serverX ~]# useradd -s /sbin/nologin -G marketing brian
```

- 4.3. Add the SMB user **brian** to Samba. The Samba user is automatically mapped to the local system user **brian**.

```
[root@serverX ~]# smbpasswd -a brian
New SMB password: redhat
Retype new SMB password: redhat
Added user brian.
```

- 4.4. Create the system user **rob** on serverX.

```
[root@serverX ~]# useradd -s /sbin/nologin rob
```

- 4.5. Add the SMB user **rob** to Samba. The Samba user is automatically mapped to the local system user **rob**.

```
[root@serverX ~]# smbpasswd -a rob
New SMB password: redhat
Retype new SMB password: redhat
Added user rob.
```

- 5. Start and enable the **smb** and **nmb** services, and allow access to them through the firewall on serverX.

- 5.1. Start the **smb** and **nmb** services on the serverX system.

```
[root@serverX ~]# systemctl start smb nmb
```

- 5.2. Enable the **smb** and **nmb** services to start at system boot on serverX.

```
[root@serverX ~]# systemctl enable smb nmb
```

- 5.3. Configure firewalld to allow access to the SMB service on serverX.

```
[root@serverX ~]# firewall-cmd --permanent --add-service=samba
success
```

```
[root@serverX ~]# firewall-cmd --reload
success
```

- 6. Verify the newly created SMB share works as expected on the desktopX system with the created Samba-only users **brian** and **rob**. The user **brian** has read and write access to the **smbshare** SMB share. The user **rob** has read access to the **smbshare** SMB share.

- 6.1. Install the *cifs-utils* package because it provides the **mount.cifs** command.

```
[root@desktopX ~]# yum -y install cifs-utils
```

- 6.2. Create the mount point **/mnt/brian** on desktopX.

```
[root@desktopX ~]# mkdir /mnt/brian
```

- 6.3. Mount the **//serverX/smbshare** Samba share temporarily as user **brian** on the mount point **/mnt/brian** on the desktopX system.

```
[root@desktopX ~]# mount -o username=brian //serverX/smbshare /mnt/brian
Password for brian@//serverX/smbshare: redhat
```

- 6.4. Verify on desktopX that user **brian** has read and write access to the **smbshare** share provided by serverX because he is a member of the marketing auxiliary group.

```
[root@desktopX ~]# echo "Hello World" >/mnt/brian/brian1.txt
[root@desktopX ~]# cat /mnt/brian/brian1.txt
Hello World
```

- 6.5. Create the mount point **/mnt/rob** on desktopX.

```
[root@desktopX ~]# mkdir /mnt/rob
```

- 6.6. Mount the **//serverX/smbshare** Samba share temporarily as user **rob** on the mount point **/mnt/rob** on the desktopX system.

```
[root@desktopX ~]# mount -o username=rob //serverX/smbshare /mnt/rob
Password for rob@//serverX/smbshare: redhat
```

- 6.7. Verify on desktopX that user **rob** has no write permission to the **smbshare** Samba share provided by the serverX system.

```
[root@desktopX ~]# touch /mnt/rob/rob1.txt
touch: cannot touch `/mnt/rob/rob1.txt': Permission denied
```

- 6.8. Test if read access for Samba user **rob** works as expected on the **/mnt/rob** Samba share on the desktopX system.

```
[root@desktopX ~]# cat /mnt/rob/brian1.txt
Hello World
```

- 7. Run the grading script on serverX:

```
[student@serverX ~]$ lab smbshare grade
```

Performing a Multiuser SMB Mount

Objectives

After completing this section, students should be able to mount an SMB share with the **multiuser** mount option, using password-based authentication and **cifscreds** to control access.

- Perform a multiuser SMB mount.
- Access the multiuser SMB mount.

Multiuser mounts with Samba

When a Samba share is mounted, the mount credentials determine the access permissions on the mount point by default. The new **multiuser** mount option separates the mount credentials from the credentials used to determine file access for each user. In Red Hat Enterprise Linux 7, this can be used with **sec=ntlmssp** authentication (contrary to the **mount.cifs(8)** man page).

The root user mounts the share using the **multiuser** option and an SMB username that has minimal access to the contents of the share. Regular users can then stash their own SMB usernames and passwords in the current session's kernel keyring with the **cifscreds** command. Their accesses to the share are authenticated with their own credentials from the keyring, not the mount credentials. The users can clear or change their credentials for that login session at any time, and they are cleared when the session ends. File access permissions are enforced entirely by the SMB server based on the access credentials currently in use.

For example, to create a new mount point **/mnt/multiuser** and mount the share **myshare** from the SMB file server **serverX**, authenticating as SMB user **fred**, who has the NTLM password **redhat**, and using the **multiuser** mount option:

```
[root@desktopX ~]# mkdir /mnt/multiuser
[root@desktopX ~]# mount -o multiuser,sec=ntlmssp,username=fred \
> //serverX/myshare /mnt/multiuser
Password for fred@//serverX/myshare: redhat
```

The command **cifscreds** is required to store authentication credentials in the keyring of the local user. Those authentication credentials are forwarded to the Samba server on a multiuser mount. The **cifs-utils** package provides the **cifscreds** command, so it is required on the desktopX system.

```
[root@desktopX ~]# yum -y install cifs-utils
```

The **cifscreds** command has various actions:

- **add** to add SMB credentials to the session keyring of a user. This option is followed by the host name of the SMB file server.
- **update** to update existing credentials in the session keyring of the user. This option is followed by the host name of the SMB file server.
- **clear** to remove a particular entry from the session keyring of the user. This option is followed by the host name of the Samba server.

- **clearall** to clear all existing credentials from the session keyring of the user.



Note

By default, **cifscreds** assumes that the username to use with the SMB credentials matches the current Linux username. A different username can be used for SMB credentials with the **-u *username*** option after the **add**, **update**, or **clear** action.

For example, assume that root has mounted **//serverX/myshare** on the mount point **/mnt/multiuser** using the **multiuser** option. In order to access files on that share, user **frank** must use **cifscreds** to temporarily stash his username and password in the kernel-managed session keyring.

```
[frank@desktopX ~]$ cifscreds add serverX
Password: redhat
[frank@desktopX ~]$ echo "Frank was here" >/mnt/multiuser/frank2.txt
[frank@desktopX ~]$ cat /mnt/multiuser/frank2.txt
Frank was here.
[frank@desktopX ~]$ exit
```

Assume that the permissions on the files in the SMB share grant **frank** read-write access to the directory, but **jane** is only granted read access.

```
[jane@desktopX ~]$ cifscreds add serverX
Password: redhat
[jane@desktopX ~]$ echo "Jane was not here" >/mnt/multiuser/jane2.txt
-bash: /mnt/multiuser/jane2.txt: Permission denied
[jane@desktopX ~]$ cat /mnt/multiuser/frank2.txt
Frank was here
```



References

mount(8), **mount.cifs(8)**, and **cifscreds(1)** man pages

Practice: Performing a Multiuser SMB Mount

Guided exercise

In this lab, you will mount a SMB share with the multiuser option.

Resources:

Machines:	desktopX and serverX
------------------	----------------------

Outcomes:

The SMB share provided by serverX is mounted on the desktopX system with the multiuser mount option.

Before you begin...

- Reset the serverX system.
- Reset the desktopX system.
- Log into and set up your server system.

```
[student@serverX ~]$ lab smbmultiuser setup
```

- Log into and set up your desktop system.

```
[student@desktopX ~]$ lab smbmultiuser setup
```

- Become **root** on your **serverX** system.

```
[student@serverX ~]$ sudo -i
```

- Become **root** on your **desktopX** system.

```
[student@desktopX ~]$ sudo -i
```

Mount the SMB share **//serverX/smbshare** permanently on the desktopX system according to the following requirements:

- The mount point on the desktopX system is the newly created directory **/mnt/multiuser**.
- The SMB share is mounted with a newly created credentials file **/root/smb-multiuser.txt**. The credentials used to mount the SMB share are username **brian** and password **redhat**.
- The SMB share is mounted with the multiuser mount option enabled.
- The already existing user **brian** on the desktopX system has a corresponding SMB account on serverX. Associate the system user **brian** on the desktopX system with the SMB user **brian** on the serverX system to access the **/mnt/multiuser** mount point. The password for **brian** is **redhat**. Verify that user **brian** has read and write access to the mounted SMB share.

- The already existing user **rob** on the desktopX system has a corresponding SMB account on serverX. Associate the system user **rob** on the desktopX system with the SMB user **rob** on the serverX system to access the **/mnt/multiuser** mount point. The password for **rob** is **redhat**. Verify that user **rob** has read but no write access to the mounted SMB share.

- 1. Install the *cifs-utils* RPM package on the desktopX system because it contains the **cifscreds** command required to store and forward authentication credentials to the Samba server with a multiuser mount.

```
[root@desktopX ~]# yum -y install cifs-utils
```

- 2. Mount the Samba share permanently on the **/mnt/multiuser** mount point on desktopX and authenticate with a credentials file. Mount the Samba share with the credentials of user **brian**.

- 2.1. Create the mount point **/mnt/multiuser** on desktopX.

```
[root@desktopX ~]# mkdir /mnt/multiuser
```

- 2.2. Create the credentials file **/root/smb-multiuser.txt** with the username and password of user **brian** on the desktopX system.

```
[root@desktopX ~]# echo 'username=brian' >/root/smb-multiuser.txt  
[root@desktopX ~]# echo 'password=redhat' >>/root/smb-multiuser.txt
```

- 2.3. Create the entry in **/etc/fstab** to permanently mount the Samba share with multiuser option as user **brian** on the **/mnt/multiuser** mount point on desktopX.

```
//serverX/smbshare /mnt/multiuser cifs credentials=/root/smb-  
multiuser.txt,multiuser,sec=ntlmssp 0 0
```

- 2.4. Verify the entry in **/etc/fstab** to permanently mount the Samba share on desktopX is correct by mounting the share with the **fstab** entry.

```
[root@desktopX ~]# mount /mnt/multiuser
```

- 3. Access the SMB multiuser mount **/mnt/multiuser** as the already existing user **brian** on desktopX. Automatically authenticate to Samba as the corresponding Samba user that exists with the same name on serverX and has read and write permission on the SMB share.

- 3.1. Switch to user **brian** on the terminal on desktopX.

```
[root@desktopX ~]# su - brian
```

- 3.2. Try to write to the **/mnt/multiuser** mount point with user **brian** on the desktopX system.

```
[brian@desktopX ~]$ touch /mnt/multiuser/testfile.txt
touch: cannot touch `test.txt`: Permission denied
```

- 3.3. Record the Samba credentials for the local user **brian** with the **cifscreds** command on desktopX.

```
[brian@desktopX ~]$ cifscreds add serverX
Password: redhat
```

- 3.4. Verify user **brian** has read and write permission on the mount point **/mnt/multiuser** on desktopX.

```
[brian@desktopX ~]$ echo "Multiuser" >/mnt/multiuser/brian2.txt
[brian@desktopX ~]$ cat /mnt/multiuser/brian2.txt
Multiuser
```

- 3.5. Exit the shell on the desktopX system.

```
[brian@desktopX ~]$ exit
[root@desktopX ~]#
```

- 4. Access the SMB multiuser mount **/mnt/multiuser** as the already existing user **rob** on desktopX. Automatically authenticate to Samba as the corresponding Samba user **rob** that exists on serverX with the same name and has read permission on the SMB share.

- 4.1. Switch to user **rob** on the terminal.

```
[root@desktopX ~]# su - rob
```

- 4.2. Record the Samba credentials for the local user **rob** with **cifscreds** on desktopX.

```
[rob@desktopX ~]$ cifscreds add serverX
Password: redhat
```

- 4.3. Verify user **rob** has read but no write permission on the mount point **/mnt/multiuser** on desktopX.

```
[rob@desktopX ~]$ echo "Multiuser" >/mnt/multiuser/rob2.txt
-bash: /mnt/multiuser/rob2.txt: Permission denied
[rob@desktopX ~]$ cat /mnt/multiuser/brian2.txt
Multiuser
```

- 5. Open a new terminal on desktopX and run the grading script:

```
[student@desktopX ~]$ lab smbmultiuser grade
```


Lab: Providing File-based Storage

Performance checklist

In this lab, you will create storage shares with NFS and SMB.

Resources:**Machines:**

desktopX and serverX

Outcomes:

Share directories with NFS and Kerberos security on serverX according to the requirements given, then mount them on desktopX; share a directory with SMB on serverX according to the requirements given, then mount it on desktopX.

Before you begin...

- Reset the serverX system.
- Reset the desktopX system.
- Log into and set up your server system.

```
[student@serverX ~]$ lab storageshares setup
```

- Log into and set up your desktop system.

```
[student@desktopX ~]$ lab storageshares setup
```

- Become **root** on your **serverX** system.

```
[student@serverX ~]$ sudo -i
```

- Become **root** on your **desktopX** system.

```
[student@desktopX ~]$ sudo -i
```

- Successfully grade the setup on desktopX and serverX with **lab storageshares grade**.

Configure the NFS server on serverX to meet the following requirements:

- Share the newly created **/krbnfs** directory on serverX with krb5p security.
- Allow read and write access on the share from the desktopX system.
- SELinux labels are exported.
- Preconfigured krb5 keytabs for the serverX and desktopX systems are available at:

<http://classroom.example.com/pub/keytabs/serverX.keytab>.

<http://classroom.example.com/pub/keytabs/desktopX.keytab>.

- Allow access to the NFS service through the firewall.

Mount the NFS share **serverX:/krbnfs** on the newly created mount point **/mnt/secureospace** permanently on desktopX with **krb5p** security and exported SELinux labels.

Configure a SMB share on the serverX system according to the following requirements:

- Share the newly created directory **/sambaspace** with SMB.
- Members of the auxiliary group **sales** have read and write permissions on the share.
- All users that are not members of the **sales** group have read permission.
- The Samba server is in the **salesdep** workgroup, and the share name in Samba is **smbspace**.
- Create the Samba-only user **frank**, who is part of the marketing team, with the password **redhat**.
- Create the new Samba-only user **martin**, who is not part of the marketing team, with the password **redhat**.
- Allow access to the SMB service through the firewall.

Mount the Samba share permanently on the **/mnt/salesshare** mount point on desktopX as a multiuser mount. Authenticate with the newly created credentials file **/root/smb-multiuser.txt**. Mount the Samba share with the credentials of user **frank**.

1. Configure NFS to share the newly created **/krbnfs** directory on serverX with **krb5p** security. Allow read and write access from the desktopX system. Mount the NFS share at the **/mnt/secureospace** directory on the desktopX system. The SELinux labels from the exported directory are available on the desktopX system where the share is mounted. The Kerberos-authenticated user **ldapuserX** can successfully access the share.
2. Mount the NFS share **serverX:/krbnfs** on the newly created mount point **/mnt/secureospace** permanently on desktopX with **krb5p** security and exported SELinux labels.
3. Configure the serverX machine to share the newly created directory **/sambaspace** with SMB. Members of the auxiliary group **sales** have read and write permissions on the share. All users that are not member of the **sales** group have read permission. Ensure that all newly created files of a user that is a member of the **sales** group have the group ownership **sales**. The Samba server is in the **salesdep** workgroup and the share name in Samba is **smbspace**.
4. Create the Samba-only user **frank**, who is part of the sales group. The user **frank** has read and write access to the **smbspace** SMB share. A new Samba user **martin** is created, who is not part of the sales team. The user **martin** has read access to the **smbspace** SMB share. Both newly added users have the SMB password **redhat**.
5. Mount the Samba share permanently on the **/mnt/salesshare** mount point on desktopX as a multiuser mount. Authenticate with the newly created credentials file **/root/smb-multiuser.txt**. Mount the Samba share with the credentials of user **frank**.
6. Successfully grade the setup on desktopX and serverX with **lab storageshares grade**.

```
[root@desktopX ~]# lab storageshares grade  
[root@serverX ~]# lab storageshares grade
```

Solution

In this lab, you will create storage shares with NFS and SMB.

Resources:

Machines:	desktopX and serverX
------------------	----------------------

Outcomes:

Share directories with NFS and Kerberos security on serverX according to the requirements given, then mount them on desktopX; share a directory with SMB on serverX according to the requirements given, then mount it on desktopX.

Before you begin...

- Reset the serverX system.
- Reset the desktopX system.
- Log into and set up your server system.

```
[student@serverX ~]$ lab storageshares setup
```

- Log into and set up your desktop system.

```
[student@desktopX ~]$ lab storageshares setup
```

- Become **root** on your **serverX** system.

```
[student@serverX ~]$ sudo -i
```

- Become **root** on your **desktopX** system.

```
[student@desktopX ~]$ sudo -i
```

- Successfully grade the setup on desktopX and serverX with **lab storageshares grade**.

Configure the NFS server on serverX to meet the following requirements:

- Share the newly created **/krbnfs** directory on serverX with krb5p security.
- Allow read and write access on the share from the desktopX system.
- SELinux labels are exported.
- Preconfigured krb5 keytabs for the serverX and desktopX systems are available at:
<http://classroom.example.com/pub/keytabs/serverX.keytab>.
<http://classroom.example.com/pub/keytabs/desktopX.keytab>.
- Allow access to the NFS service through the firewall.

Mount the NFS share **serverX:/krbnfs** on the newly created mount point **/mnt/secureospace** permanently on desktopX with **krb5p** security and exported SELinux labels.

Configure a SMB share on the serverX system according to the following requirements:

- Share the newly created directory **/sambaspace** with SMB.
- Members of the auxiliary group **sales** have read and write permissions on the share.
- All users that are not members of the **sales** group have read permission.
- The Samba server is in the **salesdep** workgroup, and the share name in Samba is **smbspace**.
- Create the Samba-only user **frank**, who is part of the marketing team, with the password **redhat**.
- Create the new Samba-only user **martin**, who is not part of the marketing team, with the password **redhat**.
- Allow access to the SMB service through the firewall.

Mount the Samba share permanently on the **/mnt/salesshare** mount point on desktopX as a multiuser mount. Authenticate with the newly created credentials file **/root/smb-multiuser.txt**. Mount the Samba share with the credentials of user **frank**.

1. Configure NFS to share the newly created **/krbnfs** directory on serverX with **krb5p** security. Allow read and write access from the desktopX system. Mount the NFS share at the **/mnt/secureospace** directory on the desktopX system. The SELinux labels from the exported directory are available on the desktopX system where the share is mounted. The Kerberos-authenticated user **ldapuserX** can successfully access the share.

- 1.1. Install the keytab provided at <http://classroom.example.com/pub/keytabs/serverX.keytab> on the serverX system.

```
[root@serverX ~]# wget -O /etc/krb5.keytab http://classroom.example.com/pub/keytabs/serverX.keytab
```

- 1.2. Enable NFS version 4.2 on the serverX system to export the SELinux labels. To do that, change the **RPCNFSDARGS=""** line in the **/etc/sysconfig/nfs** file to:

```
RPCNFSDARGS="-V 4.2"
```

- 1.3. Start the **nfs-secure-server** service on the serverX system.

```
[root@serverX ~]# systemctl start nfs-secure-server
```

- 1.4. Enable the **nfs-secure-server** to start at system boot on serverX.

```
[root@serverX ~]# systemctl enable nfs-secure-server
```

- 1.5. Create the directory **/krbnfs** on serverX.

```
[root@serverX ~]# mkdir /krbnfs
```

- 1.6. Add the directory **/krbnfs** to the **/etc/exports** file to export it with NFS. Enable krb5p security to secure access to the NFS share. Allow read and write access to the exported directory.

```
[root@serverX ~]# echo '/krbnfs desktopX(sec=krb5p,rw)' >>/etc/exports
```

- 1.7. Reload the **/etc/exports** file on the serverX system.

```
[root@serverX ~]# exportfs -r
```

- 1.8. Configure firewalld to allow access to the NFS service on serverX.

```
[root@serverX ~]# firewall-cmd --permanent --add-service=nfs
```

- 1.9. Reload the **firewalld** configuration to allow access to the NFS service on serverX.

```
[root@serverX ~]# firewall-cmd --reload
```

2. Mount the NFS share **serverX:/krbnfs** on the newly created mount point **/mnt/secureospace** permanently on desktopX with **krb5p** security and exported SELinux labels.

- 2.1. Install the keytab provided at <http://classroom.example.com/pub/keytabs/desktopX.keytab> on the desktopX system.

```
[root@desktopX ~]# wget -O /etc/krb5.keytab http://classroom.example.com/pub/keytabs/desktopX.keytab
```

- 2.2. Start the **nfs-secure** service on desktopX to help with negotiating authentication with a Kerberized NFS share.

```
[root@desktopX ~]# systemctl start nfs-secure
```

- 2.3. Enable the **nfs-secure** share to start at system boot on desktopX.

```
[root@desktopX ~]# systemctl enable nfs-secure
```

- 2.4. Create the mount point **/mnt/secureospace** on the desktopX system.

```
[root@desktopX ~]# mkdir /mnt/secureospace
```

- 2.5. Create the entry in the `/etc/fstab` file to mount the `/krbnfs` share exported by the `serverX` system on the `/mnt/secureospace` mount point on `desktopX` so that the SELinux labels from the share are shown on the mount point.

```
serverX:/krbnfs /mnt/secureospace nfs defaults,v4.2,sec=krb5p 0 0
```

- 2.6. Mount the exported NFS share on the newly created `/mnt/secureospace` directory on the `desktopX` system and verify the `/etc/fstab` entry works as expected.

```
[root@desktopX ~]# mount /mnt/secureospace
```

3. Configure the `serverX` machine to share the newly created directory `/sambaspace` with SMB. Members of the auxiliary group `sales` have read and write permissions on the share. All users that are not member of the `sales` group have read permission. Ensure that all newly created files of a user that is a member of the `sales` group have the group ownership `sales`. The Samba server is in the `salesdep` workgroup and the share name in Samba is `smbospace`.

- 3.1. Deploy the required RPM packages to run the SMB service on `serverX`.

```
[root@serverX ~]# yum -y install samba
```

- 3.2. Create the auxiliary system group `sales` on the `serverX` system.

```
[root@serverX ~]# groupadd -r sales
```

- 3.3. Create the `/sambaspace` directory on `serverX`.

```
[root@serverX ~]# mkdir -p /sambaspace
```

- 3.4. Change the group ownership of the `/sambaspace` directory to `sales` on the `serverX` system.

```
[root@serverX ~]# chgrp sales /sambaspace
```

- 3.5. Adjust the permissions on the `/sambaspace` directory to have the SGID bit set and write is prohibited by others.

```
[root@serverX ~]# chmod 2775 /sambaspace
```

- 3.6. Add the directory `/sambaspace` and all files shown as follows to the SELinux policy as a directory sharing files with SMB by setting its label to `samba_share_t`.

```
[root@serverX ~]# semanage fcontext -a -t samba_share_t '/sambaspace(/.*)?'
```

- 3.7. Apply the SELinux rule for the `/sambaspace` directory that was added in the previous step on `serverX`.

```
[root@serverX ~]# restorecon -vvFR /sambaspace
restorecon reset /sambaspace context unconfined_u:object_r:default_t:s0-
>system_u:object_r:samba_share_t:s0
```

- 3.8. Change the `/etc/samba/smb.conf` configuration file on serverX. The workgroup is **salesdep**. The newly created SMB share is named **smbspace** and shares the `/sambaspace` directory where all members of the system group **sales** have write access.

Modify or confirm the following:

```
[global]
...
workgroup = salesdep
...
security = user
passdb backend = tdbsam
```

Add a section at the end of the file as follows.

```
[smbspace]
path = /sambaspace
write list = @sales
```

4. Create the Samba-only user **frank**, who is part of the sales group. The user **frank** has read and write access to the **smbspace** SMB share. A new Samba user **martin** is created, who is not part of the sales team. The user **martin** has read access to the **smbspace** SMB share. Both newly added users have the SMB password **redhat**.

- 4.1. Install the `samba-client` RPM package because it contains the **smbpasswd**.

```
[root@serverX ~]# yum -y install samba-client
```

- 4.2. Create the system user **frank** as a member of the auxiliary group **sales** on serverX.

```
[root@serverX ~]# useradd -s /sbin/nologin -G sales frank
```

- 4.3. Add the SMB user **frank** to Samba. The Samba user is automatically mapped to the local system user **frank**.

```
[root@serverX ~]# smbpasswd -a frank
New SMB password: redhat
Retype new SMB password: redhat
Added user frank.
```

- 4.4. Create the system user **martin** on serverX.

```
[root@serverX ~]# useradd -s /sbin/nologin martin
```


- 4.5. Add the SMB user **martin** to Samba. The Samba user is automatically mapped to the local system user **martin**.

```
[root@serverX ~]# smbpasswd -a martin
New SMB password: redhat
Retype new SMB password: redhat
Added user martin.
```

- 4.6. Start the **smb** and **nmb** services on the serverX system.

```
[root@serverX ~]# systemctl start smb nmb
```

- 4.7. Enable the **smb** and **nmb** services to start at system boot on serverX.

```
[root@serverX ~]# systemctl enable smb nmb
```

- 4.8. Configure firewalld to allow access to the SMB service on serverX.

```
[root@serverX ~]# firewall-cmd --permanent --add-service=samba
success
```

```
[root@serverX ~]# firewall-cmd --reload
success
```

5. Mount the Samba share permanently on the **/mnt/salesshare** mount point on desktopX as a multiuser mount. Authenticate with the newly created credentials file **/root/smb-multiuser.txt**. Mount the Samba share with the credentials of user **frank**.

- 5.1. Install the *cifs-utils* package to provide the **mount.cifs** command on the desktopX system.

```
[root@desktopX ~]# yum -y install cifs-utils
```

- 5.2. Create the mount point **/mnt/salesshare** on desktopX.

```
[root@desktopX ~]# mkdir /mnt/salesshare
```

- 5.3. Mount the Samba share permanently with a credentials file on the **/mnt/salesshare** mount point on desktopX. Mount the Samba share with the credentials of user **frank**.

```
[root@desktopX ~]# echo 'username=frank' >/root/smb-multiuser.txt
[root@desktopX ~]# echo 'password=redhat' >>/root/smb-multiuser.txt
```

- 5.4. Create the entry in **/etc/fstab** to permanently mount the Samba share **//serverX/smbspace** with multiuser option as user frank on the **/mnt/salesshare** mount point on desktopX.

```
//serverX/smbspace /mnt/salesshare cifs credentials=/root/smb-  
multiuser.txt,multiuser,sec=ntlmssp 0 0
```

- 5.5. Verify the entry in **/etc/fstab** to permanently mount the Samba share on desktopX is correct by mounting the share with the fstab entry.

```
[root@desktopX ~]# mount /mnt/salesshare
```

6. Successfully grade the setup on desktopX and serverX with **lab storageshares grade**.

```
[root@desktopX ~]# lab storageshares grade  
[root@serverX ~]# lab storageshares grade
```

Summary

Exporting NFS File Systems

In this section, students learned how to:

- Describe NFS server configuration.
- Export an IP-based NFS share.

Protecting NFS Exports

In this section, students learned how to:

- Set up Kerberos-enabled NFS.
- Configure labeled NFS and SELinux.

Providing SMB File Shares

In this section, students learned how to perform SMB server configuration.

Performing a Multiuser SMB Mount

In this section, students learned how to perform a multiuser SMB mount.



CHAPTER 9

CONFIGURING MARIADB DATABASES

Overview	
Goal	To provide a MariaDB SQL database for use by programs and database administrators.
Objectives	<ul style="list-style-type: none">• Install MariaDB.• Configure and administer MariaDB.• Configure user and access rights.• Back up and restore MariaDB databases.
Sections	<ul style="list-style-type: none">• Installing MariaDB (and Practice)• Working with MariaDB Databases (and Practice)• Managing Database Users and Access Rights (and Practice)• Creating and Restoring MariaDB Backups (and Practice)
Lab	<ul style="list-style-type: none">• Configuring MariaDB Databases

Installing MariaDB

Objectives

After completing this section, students should be able to describe a relational database and be able to install and perform a basic configuration of a simple MariaDB relational database server.

Relational databases

A relational database is a mechanism that allows the persistence of data in an organized way. Databases store data items organized as a set of tables, with each table representing an entity. In a given table, each row corresponds to a record, while each column corresponds to an attribute of that record.

```
MariaDB [inventory]> SELECT * FROM product;
+-----+-----+-----+-----+-----+-----+
| id | name           | price  | stock | id_category | id_manufacturer |
+-----+-----+-----+-----+-----+-----+
| 1 | ThinkServer TS140 | 539.88 | 20   | 2           | 4               |
| 2 | ThinkServer TS440 | 1736.00 | 10   | 2           | 4               |
| 3 | RT-AC68U         | 219.99 | 10   | 1           | 3               |
| 4 | X110 64GB        | 73.84  | 100  | 3           | 1               |
+-----+-----+-----+-----+-----+-----+
4 rows in set (0.00 sec)
```

```
MariaDB [inventory]> SELECT * FROM category;
+-----+-----+
| id | name      |
+-----+-----+
| 1 | Networking |
| 2 | Servers   |
| 3 | Ssd       |
+-----+-----+
3 rows in set (0.00 sec)
```

```
MariaDB [inventory]> SELECT * FROM manufacturer;
+-----+-----+-----+-----+
| id | name      | seller           | phone_number |
+-----+-----+-----+-----+
| 1 | SanDisk   | John Miller     | +1 (941) 329-8855 |
| 2 | Kingston  | Mike Taylor     | +1 (341) 375-9999 |
| 3 | Asus      | Wilson Jackson  | +1 (432) 367-8899 |
| 4 | Lenovo    | Allen Scott     | +1 (876) 213-4439 |
+-----+-----+-----+-----+
4 rows in set (0.00 sec)
```

The previous tables show:

- The **product** table has four records. Each record has six attributes: (**id**, **name**, **price**, **stock**, **id_category**, **id_manufacturer**).
- **X110 64GB** is an SSD manufactured by SanDisk.
- The **seller** responsible for the ThinkServer TS140 product is Allen Scott.

There are two relational database packages provided with Red Hat Enterprise Linux 7:

- **PostgreSQL**—An open source database developed by the PostgreSQL Global Development Group, consisting of Postgres users (both individuals and companies) and other companies and volunteers, supervised by companies such as Red Hat and EnterpriseDB.
- **MariaDB**—A community-developed branch of MySQL built by some of the original authors of MySQL. It offers a rich set of feature enhancements, including alternate storage engines, server optimizations, and patches. The MariaDB Foundation works closely and cooperatively with the larger community of users and developers in the spirit of free and open source software.



Note

MariaDB is a community-developed replacement for MySQL. MySQL is not shipped in the core toolset for Red Hat Enterprise Linux 7, but it is available through Red Hat Software Collections 1.1.

Red Hat Software Collections is a fully-supported parallel set of tools included as part of most RHEL 7 subscriptions. Red Hat Software Collections provides a set of dynamic programming languages, database servers, web servers and various related packages that are either more recent than the version shipped in the core RHEL distribution or are not otherwise available in it at all. Packages in Red Hat Software Collections have a faster but shorter support cycle (up to 3 years of support).

More information is available in the *Red Hat Software Collections 1.1 Release Notes* at <https://access.redhat.com/docs>.

MariaDB installation

A full MariaDB database installation requires both the *mariadb* and *mariadb-client* groups of software to be installed.

The following packages will be installed with the **mariadb** group:

- *mariadb-server*—The MariaDB server and related files (mandatory package).
- *mariadb-bench*—MariaDB benchmark scripts and data (optional package).
- *mariadb-test*—The test suite distributed with MariaDB (optional package).

The following packages will be installed with the **mariadb-client** group:

- *mariadb*—A community-developed branch of MySQL (mandatory package).
- *MySQL-python*—A MariaDB interface for Python (default package).
- *mysql-connector-odbc*—ODBC driver for MariaDB (default package).
- *libdbi-dbd-mysql*—MariaDB plug-in for libdbi (optional package).
- *mysql-connector-java*—Native Java driver for MariaDB (optional package).
- *perl-DBD-MySQL*—A MariaDB interface for Perl (optional package).

The `/etc/my.cnf` file has default configurations for MariaDB, such as the data directory, socket bindings, and log and error file location.



Note

Instead of adding new configurations to the `/etc/my.cnf` file, a newly created file named `*.cnf` can be added to the `/etc/my.cnf.d/` directory holding the configuration of MariaDB.

MariaDB installation demonstration

In this example, please follow along with these steps while your instructor demonstrates how to install a MariaDB database server.

1. Install MariaDB on serverX with the `yum` command.

```
[root@serverX ~]# yum groupinstall mariadb mariadb-client -y
```

2. Start the MariaDB service on serverX with the `systemctl` command.

```
[root@serverX ~]# systemctl start mariadb
```



Note

The default MariaDB log file is `/var/log/mariadb/mariadb.log`. This file should be the first place to look when troubleshooting MariaDB.

3. Enable the MariaDB service to start at boot on serverX.

```
[root@serverX ~]# systemctl enable mariadb
```

4. Verify the status of the service on serverX.

```
[root@serverX ~]# systemctl status mariadb
```

The `status` option reports some attributes, if the database is started:

- Loaded—Shows if this service is loaded and enabled.
- Active—Shows if this service is activated.
- Main PID—Shows the main process ID from this service.
- CGroup—Shows all processes that belong to this service.



Note

If the database is stopped, the **status** option will report the last known PID, and that the service is inactive.

Improve MariaDB installation security

MariaDB provides a program to improve security from the baseline install state. Run **mysql_secure_installation** without arguments:

```
[root@serverX ~]# mysql_secure_installation
```

This program enables improvement of MariaDB security in the following ways:

- Sets a password for **root** accounts.
- Removes **root** accounts that are accessible from outside the local host.
- Removes **anonymous-user** accounts.
- Removes the **test** database.

The script is fully interactive, and will prompt for each step in the process.

MariaDB and networking

MariaDB can be configured to be accessed remotely, or limited to just local connections.

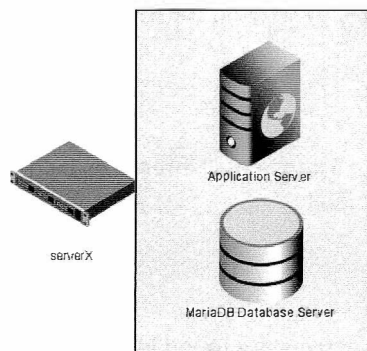


Figure 9.1: Local access to MariaDB




Figure 9.2: Remote access to MariaDB

In the first scenario, the database can only be accessed locally. Security is greatly improved, because the records will only be accessed from applications that are on the same server. The disadvantage is that the server will share the same resources with other services, and this may impact performance in the database server.

In the second scenario, the database can be accessed remotely. In this case, safety decreases because another port is opened on the server, which may result in an attack. On the other hand, the performance of the server increases by not having to share resources.

When MariaDB is accessed remotely, by default, the server listens for TCP/IP connections on all available interfaces on port 3306.



Note

Although the MariaDB service listens on all interfaces by default, no users have remote access permission, also by default.

Configuring MariaDB networking

MariaDB network configuration directives are found in the `/etc/my.cnf` file, under the `[mysqld]` section.

bind-address

The server will listen based on this directive. Only one value may be entered. Possible values are:

- Host name
- IPv4 address
- IPv6 address
- Set this value to `:::` to connect to all available addresses (IPv6 and IPv4), or leave blank (or set to `0.0.0.0`) for all IPv4 addresses.



Important

There can be only one **bind-address** entry in `/etc/my.cnf`. On a system with multiple addresses, selecting a single address is possible, or all addresses, but nothing in between.

skip-networking

If set to 1, the server will listen only for local clients. All interaction with the server will be through a socket, located by default at `/var/lib/mysql/mysql.sock`. This location can be changed with a **socket** value in `/etc/my.cnf`.



Important

Be aware that if networking is shut off in this manner, this disables connections via localhost as well. The MySQL client can still make local connections through the socket file automatically.

port

Port to listen on for TCP/IP connections.



Important

For remote access, the firewall needs to be modified. Fortunately, it is a known service, so it can be simply added via:

```
[root@serverX ~]# firewall-cmd --permanent --add-service=mysql
[root@serverX ~]# firewall-cmd --reload
```

R

References

`mysql_secure_installation(1)`, `mysql.server(1)`, and `mysqld_selinux(8)` man pages

Red Hat Software Collections documentation
https://access.redhat.com/site/documentation/en-US/Red_Hat_Software_Collections/

Practice: Installing MariaDB

Guided exercise

In this lab, you will install a MariaDB database server.

Resources:	
Machines:	desktopX and serverX

Outcomes:

A MariaDB database server running on your serverX machine.

Before you begin...

- Reset your **desktopX** machine.
- Reset your **serverX** machine.

You have been asked to install a MariaDB database server on your serverX machine. You need to secure the MariaDB service, and must configure it to accept connections only from local clients.

- 1. Install the *mariadb* and *mariadb-client* groups.

□ 1.1.

```
[root@serverX ~]# yum groupinstall mariadb mariadb-client -y
```

- 2. Start and enable the **mariadb** service.

□ 2.1.

```
[root@serverX ~]# systemctl start mariadb
```

□ 2.2.

```
[root@serverX ~]# systemctl enable mariadb
```

- 3. Verify that MariaDB is listening on all interfaces.

□ 3.1.

```
[root@serverX ~]# ss -tulpn | grep mysql

tcp LISTEN 0 50 *:3306 :* users:(("mysqld", 13611, 13))
```



Note

*Why grep on **mysql** when you are running MariaDB?* MariaDB is a fork of MySQL which is meant to act as a drop-in replacement.

- 4. Enable the **skip-networking** directive.

- 4.1. Open the file `/etc/my.cnf` in a text editor, and in section **[mysqld]**, add the line

```
skip-networking=1
```

- 5. Restart the **mariadb** service.

- 5.1.

```
[root@serverX ~]# systemctl restart mariadb
```

- 6. Verify that MariaDB is **not** listening on all interfaces.

- 6.1.

```
[root@serverX ~]# ss -tulpn | grep mysql
```

This command should now return nothing.

- 7. Secure the Mariadb service using the **mysql_secure_installation** tool. Set the **root** password to **redhat**, and answer Yes to all other questions.

□ 7.1.

```
[root@serverX ~]# mysql_secure_installation

In order to log into MariaDB to secure it, we'll need the current
password for the root user.  If you've just installed MariaDB, and
you haven't set the root password yet, the password will be blank,
so you should just press enter here.

Enter current password for root (enter for none):
OK, successfully used password, moving on...

Setting the root password ensures that nobody can log into the MariaDB
root user without the proper authorisation.

Set root password? [Y/n] Y
New password: redhat
Re-enter new password: redhat
Password updated successfully!
Reloading privilege tables..
... Success!

By default, a MariaDB installation has an anonymous user, allowing anyone
to log into MariaDB without having to have a user account created for
them.  This is intended only for testing, and to make the installation
go a bit smoother.  You should remove them before moving into a
production environment.

Remove anonymous users? [Y/n] Y
... Success!

Normally, root should only be allowed to connect from 'localhost'.  This
ensures that someone cannot guess at the root password from the network.

Disallow root login remotely? [Y/n] Y
... Success!

By default, MariaDB comes with a database named 'test' that anyone can
access.  This is also intended only for testing, and should be removed
before moving into a production environment.

Remove test database and access to it? [Y/n] Y
- Dropping test database...
... Success!
- Removing privileges on test database...
... Success!

Reloading the privilege tables will ensure that all changes made so far
will take effect immediately.

Reload privilege tables now? [Y/n] Y
... Success!

Cleaning up...

All done!  If you've completed all of the above steps, your MariaDB
installation should now be secure.

Thanks for using MariaDB!
```

□ 7.2. Verify that the **root** user cannot log in without a password.

```
[root@serverX ~]# mysql -u root
ERROR 1045 (28000): Access denied for user 'root'@'localhost' (using
password: NO)
```

- 7.3. Verify that the **test** database is removed.

```
[root@serverX ~]# mysql -u root -p redhat
MariaDB [(none)]> show databases;
+-----+
| Database          |
+-----+
| information_schema |
| mysql             |
| performance_schema |
+-----+
3 rows in set (0.00 sec)
MariaDB [(none)]> exit;
Bye
```

Working with MariaDB Databases

Objectives

After completing this section, students should be able to examine, search, create, and change database information using Structured Query Language (SQL) and MariaDB statements.

Creating a database

The installation of the *mariadb-client* group provides a program called **mysql**. With this program, it is possible to connect to a local or remote MariaDB database server.

```
[root@serverX ~]# mysql1 -u2 root3 -h4 localhost5 -p6
```

- 1 Client to connect to the MariaDB database server.
- 2 Option to specify the username for this connection.
- 3 Username for this connection.
- 4 Option to specify the host name for this connection. If not specified, the default value is localhost.
- 5 Host name for this connection.
- 6 Option to prompt for password.

A database in MariaDB is implemented as a directory. The default installation has four databases. To list the databases, run the command:

```
MariaDB [(none)]> SHOW DATABASES;  
+-----+  
| Database |  
+-----+  
| information_schema |  
| mysql |  
| performance_schema |  
| test |  
+-----+  
4 rows in set (0.01 sec)
```



Important

The only default database that can be erased is **test**.



Note

Unlike the shell, MariaDB—like most relational database systems—is not case-sensitive for simple commands. **show databases;** and **SHOW DATABASES;** work just as well as **ShOw DATAbases ;**. However, table and database names *are* case-sensitive. Many databases are set up using all lowercase for the database names, so common practice is to use uppercase for the commands to differentiate the command itself from the target of the command (more examples follow). The important part of these commands, and of most commands entered at the prompt, is to terminate the command with a **;**.

To create a new database, run the command:

```
MariaDB [(none)]> CREATE DATABASE inventory;
```

After the creation of the new database, the next step is to connect to this database so that it can be populated with tables and data:

```
MariaDB [(none)]> USE inventory;
```



Note

It is possible to switch between databases at any time with this command.

MariaDB (like all relational database systems) can have multiple tables per database. List the tables with the **SHOW TABLES;** command:


```

MariaDB [(none)]> USE mysql;
MariaDB [(none)]> SHOW TABLES;
+-----+
| Tables_in_mysql |
+-----+
| columns_priv |
| db |
| event |
| func |
| general_log |
| help_category |
| help_keyword |
| help_relation |
| help_topic |
| host |
| ndb_binlog_index |
| plugin |
| proc |
| procs_priv |
| proxies_priv |
| servers |
| slow_log |
| tables_priv |
| time_zone |
| time_zone_leap_second |
| time_zone_name |
| time_zone_transition |
| time_zone_transition_type |
| user |
+-----+
24 rows in set (0.00 sec)
    
```

To list attributes (or the column names) from a table, use:

```

MariaDB [(mysql)]> DESCRIBE servers;
+-----+-----+-----+-----+-----+-----+
| Field | Type | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| Server_name | char(64) | NO | PRI | | |
| Host | char(64) | NO | | | |
| Db | char(64) | NO | | | |
| Username | char(64) | NO | | | |
| Password | char(64) | NO | | | |
| Port | int(4) | NO | | 0 | |
| Socket | char(64) | NO | | | |
| Wrapper | char(64) | NO | | | |
| Owner | char(64) | NO | | | |
+-----+-----+-----+-----+-----+-----+
42 rows in set (0.00 sec)
    
```

This output completely describes the data in the **servers** table in the **mysql** database. The **Port** attribute is stored as an **integer**, using a maximum of **4** digits, and defaults to **0**.

The **Key** value is null for most of these attributes. Only the **Server_name** has a value: **PRI**. This sets the attribute as the primary key for the table. Primary keys are unique identifiers for the data in the table. No two entries can have the same primary key, and only one attribute may be set as the primary key. Primary keys are often used to link tables together, and are an important concept when designing complex databases. There are also secondary keys, and composite keys (where multiple attributes together form the unique key). A deeper discussion of keys is beyond the scope of this course.

The **Extra** value is used to show any additional features of the attribute. This value can be complex, but a common one is **auto_increment**, which states that the value of this column will be incremented by 1 for each new entry made into the table. It is a common value for primary keys to have, as seen in later examples.

Using SQL: Structured Query Language

Structured Query Language (SQL) is a special programming language designed for managing data held in relational databases. Some common SQL commands include **insert**, **update**, **delete**, and **select**.



Note

These four basic commands are often referred to by the generic term "CRUD operations." CRUD stands for Create (**insert**), Read (**select**), Update (**update**), and Delete (**delete**).

To **insert** data into a table, the first step is to figure out the attributes of the table.

```
MariaDB [(inventory)]> DESCRIBE product;
+-----+-----+-----+-----+-----+-----+
| Field          | Type          | Null | Key | Default | Extra          |
+-----+-----+-----+-----+-----+-----+
| id             | int(11)       | NO   | PRI | NULL    | auto_increment |
| name           | varchar(100)  | NO   |     | NULL    |                |
| price          | double        | NO   |     | NULL    |                |
| stock          | int(11)       | NO   |     | NULL    |                |
| id_category    | int(11)       | NO   |     | NULL    |                |
| id_manufacturer | int(11)       | NO   |     | NULL    |                |
+-----+-----+-----+-----+-----+-----+
6 rows in set (0.00 sec)
```

In this example, all attributes are required. To insert a new product, the command will be rather long and complicated:

```
MariaDB [(inventory)]> INSERT INTO
  product❶(name,price,stock,id_category,id_manufacturer)❷ VALUES ('SDSSDP-1286-G25
  2.5',82.04,30,3,1)❸;
Query OK, 1 row affected (0.00 sec)
```

- ❶ Table name.
- ❷ Attributes (columns) that will be inserted.
- ❸ Values that will be inserted in the same order defined by attributes, using the proper data types.



Note

Note that the attribute ID was not specified, even though it is required. When inserting a new record, MariaDB will automatically assign a sequential value for that column. This is because this column is marked as **auto_increment**.

Delete a record with the **delete** statement:

```
MariaDB [(inventory)]> DELETE FROM product❶ WHERE❷ id = 1❸ ;
Query OK, 1 row affected (0.01 sec)
```

- ❶ Table name.
- ❷ Clause that imposes a condition on the command execution.
- ❸ The condition for the record to be deleted; often the primary key-value pair is used.



Warning

If the **where** clause is not specified, *all* records in the table will be erased. This is the database equivalent of running `rm -rf /`.

To update a record, use an **update** statement:

```
MariaDB [(inventory)]> UPDATE product❶ SET❷ price=89.90, stock=60 WHERE❸ id = 5❹ ;
Query OK, 1 row affected (0.01 sec)
```

- ❶ Table name.
- ❷ Define the new value from specified attributes.
- ❸ Clause that imposes a condition on the command execution.
- ❹ The condition for the record to be updated.



Warning

If the **where** clause is not specified, all records will be updated.

To read data records from the database, use the **select** statement:

```
MariaDB [(inventory)]> SELECT name,price,stock❶ FROM product❷ ;
+-----+-----+-----+
| name          | price | stock |
+-----+-----+-----+
| ThinkServer TS140 | 539.88 | 20 |
| RT-AC68U      | 219.99 | 10 |
| X110 64GB     | 73.84 | 100 |
| SDSSDP-128G-G25 2.5 | 82.04 | 30 |
+-----+-----+-----+
4 rows in set (0.00 sec)
```

- ❶ Attributes that will be selected.
- ❷ Table name.

To select all attributes, use the wild card `*`:

```

MariaDB [(inventory)]> SELECT * FROM product;
+-----+-----+-----+-----+-----+-----+
| id | name                | price | stock | id_category | id_manufacturer |
+-----+-----+-----+-----+-----+-----+
| 2 | ThinkServer TS140  | 539.88 | 20 | 2 | 4 |
| 3 | RT-AC68U           | 219.99 | 10 | 1 | 3 |
| 4 | X110 64GB          | 73.84 | 100 | 3 | 1 |
| 5 | SDSSDP-128G-G25 2.5 | 82.04 | 30 | 3 | 1 |
+-----+-----+-----+-----+-----+-----+
4 rows in set (0.00 sec)

```

Filter results with the **where** clause:

```

MariaDB [(inventory)]> SELECT * FROM product WHERE price > 100;
+-----+-----+-----+-----+-----+-----+
| id | name                | price | stock | id_category | id_manufacturer |
+-----+-----+-----+-----+-----+-----+
| 2 | ThinkServer TS140  | 539.88 | 20 | 2 | 4 |
| 3 | RT-AC68U           | 219.99 | 10 | 1 | 3 |
+-----+-----+-----+-----+-----+-----+
2 rows in set (0.00 sec)

```

Common Operators for where Clauses

Operator	Description
=	Equal
<>	Not equal. Note: In some versions of SQL, this operator may be written as !=
>	Greater than
<	Less than
>=	Greater than or equal
<=	Less than or equal
BETWEEN	Between an inclusive range
LIKE	Search for a pattern
IN	To specify multiple possible values for a column



References

mysql(1) man page

MariaDB Knowledge Base: *SQL Commands*
<https://mariadb.com/kb/en/sql-commands/>

MariaDB Knowledge Base: *Data Types*
<https://mariadb.com/kb/en/data-types/>

Practice: MariaDB Commands

Quiz

Match the following items to their counterparts in the table.

DELETE FROM table_name WHERE attribute = value;	
DESCRIBE table_name;	SELECT * FROM table_name;
SHOW DATABASES;	SHOW tables;
UPDATE table_name SET attribute=value WHERE attribute > value;	
USE database_name;	

Objective	Statement
Describe a table	
Update a record in the table	
Connect to a specific database	
List databases available in MariaDB	
List tables available in a database	

Objective	Statement
Select data records from a table	
Delete records from a table	

Solution

Match the following items to their counterparts in the table.

Objective	Statement
Describe a table	DESCRIBE table_name;
Update a record in the table	UPDATE table_name SET attribute=value WHERE attribute > value;
Connect to a specific database	USE database_name;
List databases available in MariaDB	SHOW DATABASES;
List tables available in a database	SHOW tables;
Select data records from a table	SELECT * FROM table_name;
Delete records from a table	DELETE FROM table_name WHERE attribute = value;

Managing Database Users and Access Rights

Objectives

After completing this section, students should be able to configure MariaDB users and grant or revoke database access rights.

Creating user accounts with MariaDB

By default, MariaDB handles authentication and authorization through the **user** table in the **mysql** database. This means that the root password for the database is persisted in the **user** table and not in the operating system.



Note

Recent versions of MariaDB can use PAM for authentication on Linux.

The **CREATE USER** statement creates new accounts. This statement will create a new row in the **mysql.user** table that has no privileges.



Note

information_schema and **test** database allows some privileges for all users. This is the main reason why the **test** database is often deleted.



Important

To create a new user, the connected user must have the global **CREATE USER** privilege or the **INSERT** privilege for the **mysql** database.

Account names are specified as '**user_name**'@'**host_name**'. This makes it possible to create multiple user accounts with the same name, but with different privileges according to the source host (that is, the host from which the user is connecting).

```
MariaDB [(none)]> CREATE USER mobius@localhost❶ IDENTIFIED BY 'redhat'❷;
```

- ❶ Username/host name for this account
- ❷ Password for this account

This user account can only connect from **localhost**, with the password **redhat**, and has no privileges. Passwords are encrypted in the **user** table:

```
MariaDB [mysql]> SELECT host,user,password FROM user WHERE user = 'mobius';
+-----+-----+-----+
| host      | user  | password                                     |
+-----+-----+-----+
| localhost | mobius | *84BB5DF4823DA319BBF86C99624479A198E6EEE9 |
+-----+-----+-----+
1 row in set (0.00 sec)
```


When using this account, before granting any privileges, access will be denied for almost any action:

```
[root@serverX ~]# mysql -u mobius -p
Enter password: redhat
MariaDB [(none)]> create database inventory;
ERROR 1044 (42000): Access denied for user 'mobius'@'localhost' to database 'inventory'
```



Note

If the host name is not provided, it is assumed to be "%". This means that this user can access from any source host.

Account Examples

Account	Description
mobius@'localhost'	User mobius can connect just from localhost.
mobius@'192.168.1.5'	User mobius can connect from 192.168.1.5 host.
mobius@'192.168.1.%'	User mobius can connect from any host that belongs to the network 192.168.1.0.
mobius@'%'	User mobius can connect from any host.
mobius@'2000:472:18:b51:c32:a21'	User mobius can connect from 2000:472:18:b51:c32:a21 host.

Granting and revoking privileges for user accounts

Privileges are the permissions that the user may have within MariaDB. The privileges are organized as:

- Global privileges, such as **CREATE USER** and **SHOW DATABASES**, for administration of the database server itself.
- Database privileges, such as **CREATE** for creating databases and working with databases on the server at a high level.
- Table privileges, such as the CRUD commands, for creating tables and manipulating data in the database.
- Column privileges, for granting table-like command usage, but on a particular column (generally rare).
- Other, more granular privileges, which are discussed in detail in the MariaDB documentation.

The **GRANT** statement can be used to grant privileges to accounts. The connected user must have the **GRANT OPTION** privilege (a special privilege that exists at several levels) to grant privileges. A user may only grant privileges to others that have already been granted to that user (for example, **mobius** cannot grant **SELECT** privileges on a database table unless **mobius** already has that privilege *and* the **GRANT OPTION** table privilege).

```

[root@serverX ~]# mysql -u mobius -p
Enter password: redhat
MariaDB [(none)]> use inventory;
MariaDB [(inventory)]> select * from category;
ERROR 1142 (42000): SELECT command denied to user 'mobius'@'localhost' for table
'category'
MariaDB [(inventory)]> exit
[root@serverX ~]# mysql -u root -p
Enter password: redhat
MariaDB [(none)]> use inventory;
MariaDB [(inventory)]> GRANT SELECT, UPDATE, DELETE, INSERT❶ on inventory.category❷ to
mobius@localhost❸;
Query OK, 0 rows affected (0.00 sec)
MariaDB [(inventory)]> exit
[root@serverX ~]# mysql -u mobius -p
Enter password: redhat
MariaDB [(none)]> use inventory;
MariaDB [(inventory)]> select * from category;
+----+-----+
| id | name   |
+----+-----+
|  1 | Networking |
|  2 | Servers  |
|  3 | Ssd     |
+----+-----+
3 rows in set (0.00 sec)

```

- ❶ Define the privileges to be granted (in this case, the full CRUD capabilities are being granted).
- ❷ Define which table the privilege will be granted for.
- ❸ The user to be granted privileges.

Grant Examples

Grant	Description
GRANT SELECT ON database.table TO username@hostname	Grant select privilege for a specific table in a specific database to a specific user.
GRANT SELECT ON database.* TO username@hostname	Grant select privilege for all tables in a specific database to a specific user.
GRANT SELECT ON *.* TO username@hostname	Grant select privilege for all tables in all databases to a specific user.
GRANT CREATE, ALTER, DROP ON database.* to username@hostname	Grant privilege to create, alter, and drop tables in a specific database to a specific user.
GRANT ALL PRIVILEGES ON *.* to username@hostname	Grant all available privileges for all databases to a specific user, effectively creating a superuser, similar to root.

The **REVOKE** statement allows for revoking privileges from accounts. The connected user must have the **GRANT OPTION** privilege and have the privileges that are being revoked to revoke a privilege.

```
MariaDB [(none)]> REVOKE SELECT, UPDATE, DELETE, INSERT❶ on inventory.category❷ from
mobius@localhost❸;
Query OK, 0 rows affected (0.00 sec)
```

- ❶ Define the privileges to be revoked.
- ❷ Define for which table the privilege will be revoked.
- ❸ Privilege revoked from this user.



Important

After granting or revoking a privilege, reload all privileges from the privileges tables in the mysql database.

```
MariaDB [(none)]> FLUSH PRIVILEGES;
```

In order to revoke privileges, the list of privileges granted to a user will be needed. The simple command **SHOW GRANTS FOR *username***; will provide the list of privileges for that user:

```
MariaDB [(none)]> SHOW GRANTS FOR root@localhost;
+-----+
| Grants for root@localhost |
+-----+
| GRANT ALL PRIVILEGES ON *.* TO 'root'@'localhost' WITH GRANT OPTION |
| GRANT PROXY ON ''@' TO 'root'@'localhost' WITH GRANT OPTION |
+-----+
2 rows in set (0.00 sec)
```

When a user is no longer required, it can be deleted from the database using **DROP USER *username***; . The *username* should use the same '*user*'@'*host*' format used for **CREATE USER**.



Important

If an account that is currently connected is **DROP**ped, it will not be deleted until the connection is closed. The connection will not be automatically closed.

Troubleshooting database access

Some Common DB Access Issues

Issue	Solution
User has been granted access to connect from any host, but can only connect on localhost using the mysql command (applications he or she uses cannot connect, even on localhost).	Remove the skip-networking directive from my.cnf and restart the service.
User can connect with any application on localhost, but not remotely.	Check the bind-address configuration in my.cnf to ensure the database is accessible.

Issue	Solution
	Ensure that the user table includes an entry for the user from the host he is trying to connect from.
User can connect, but cannot see any database other than information_schema and test .	Common problem when a user has just been created, as they have no privileges by default, though they can connect and use those default databases. Add grants for the database the user requires.
User can connect, but cannot create any databases.	Grant the user the global CREATE privilege (this also grants DROP privileges).
User can connect, but cannot read or write any data.	Grant the user the CRUD privileges for only the database he or she will be using.



References

mysql(1) man page

MariaDB Knowledge Base: *User Account Management*
<https://mariadb.com/kb/en/user-account-management/>

MariaDB Knowledge Base: *Account Management SQL Commands: REVOKE*
<https://mariadb.com/kb/en/revoke/>

MariaDB Knowledge Base: *Account Management SQL Commands: GRANT, "Database Privileges"*
<https://mariadb.com/kb/en/grant/#database-privileges>

Practice: Managing Users

Guided exercise

In this lab, you will install a MariaDB server.

Resources:

Machines:	desktopX and serverX
------------------	----------------------

Outcomes:

One MariaDB user with select privilege and another user with insert, update, delete, and select privileges on your serverX machine.

Before you begin...

- Reset your desktopX system.
- Reset your serverX system.
- Log into and set up your serverX system.

```
[student@desktopX ~]$ lab mariadb-users setup
```

- Log into and set up your desktopX system.

```
[student@desktopX ~]$ lab mariadb-users setup
```

You have been asked to create two MariaDB users on your serverX machine, according to the following requirements:

User	Accepts connection from host	Password	Privileges
john	localhost	john_password	Insert, update, delete, select on all tables from inventory database
steve	any host	steve_password	Select on all tables from inventory database

- 1. Connect to MariaDB with user root (blank password).

- 1.1.

```
[student@serverX ~]$ mysql -u root
```

- 2. Create the users john and steve.

- 2.1.

```
MariaDB [(none)]> CREATE USER john@localhost identified by 'john_password';
MariaDB [(none)]> CREATE USER steve@%' identified by 'steve_password';
```

- 3. Grant insert, update, delete, and select privileges to user john.

□ 3.1.

```
MariaDB [(none)]> GRANT INSERT, UPDATE, DELETE, SELECT on inventory.* to john@localhost;
```

□ 4. Grant select privilege to user steve.

□ 4.1.

```
MariaDB [(none)]> GRANT SELECT on inventory.* to steve@'%';
```

□ 5. Flush the privileges.

□ 5.1.

```
MariaDB [(none)]> FLUSH PRIVILEGES;
MariaDB [(none)]> exit;
```

□ 6. Connect with user john and verify his privileges.

□ 6.1. Connect to MariaDB.

```
[student@serverX ~]$ mysql -u john -p
```

□ 6.2. Select the inventory database.

```
MariaDB [(none)]> USE inventory;
```

□ 6.3. Verify the select privilege.

```
MariaDB [(inventory)]> SELECT * FROM category;
+----+-----+
| id | name      |
+----+-----+
|  1 | Networking|
|  2 | Servers   |
|  3 | Ssd       |
+----+-----+
3 rows in set (0.00 sec)
```

□ 6.4. Verify the insert privilege.

```
MariaDB [(inventory)]> INSERT INTO category(name) VALUES('Memory');
Query OK, 1 row affected (0.00 sec)
```

□ 6.5. Verify the update privilege.

```
MariaDB [(inventory)]> UPDATE category SET name='Solid State Drive' where
id = 3;
Query OK, 1 row affected (0.01 sec)
Rows matched: 1  Changed: 1  Warnings: 0
```

□ 6.6. Verify the delete privilege.

```
MariaDB [(inventory)]> DELETE FROM category WHERE name LIKE 'Memory';  
Query OK, 1 row affected (0.01 sec)
```

- 7. On your desktopX system, connect with user steve and verify his privileges.

- 7.1. Connect to MariaDB.

```
[student@desktopX ~]$ mysql -u steve -h serverX -p
```

- 7.2. Select the inventory database.

```
MariaDB [(none)]> USE inventory;
```

- 7.3. Verify the select privilege.

```
MariaDB [(inventory)]> SELECT * FROM category;  
+----+-----+  
| id | name          |  
+----+-----+  
| 1  | Networking    |  
| 2  | Servers       |  
| 3  | Solid State Drive |  
+----+-----+  
3 rows in set (0.00 sec)  
MariaDB [(inventory)]> exit;
```

- 7.4. Verify the insert privilege.

```
MariaDB [(inventory)]> INSERT INTO category(name) VALUES('Memory');  
ERROR 1142 (42000): INSERT command denied to user  
'steve'@'desktopX.example.com' for table 'category'
```

Creating and Restoring MariaDB Backups

Objectives

After completing this section, students should be able to back up a MariaDB database safely and restore a database backup.

Creating a backup

It is very important to back up MariaDB databases, and databases in general. The database often contains most of a company's mission-critical data (sales, clients, etc.). Performing backups enables a system administrator to recover data after several types of events:

- Operating system crash
- Power failure
- File system crash
- Hardware problem
- Security breach
- Database corruption
- Data poisoning

There are two ways to back up MariaDB:

- Logical
- Physical (raw)

Logical backups export information and records in plain text files, while *physical backups* consist of copies of files and directories that store content.

Logical backups have these characteristics:

- The database structure is retrieved by querying the database.
- Logical backups are highly portable, and can be restored to another database provider (such as Postgres) in some cases.
- Backup is slower because the server must access database information and convert it to a logical format.
- Performed while the server is online.
- Backups do not include log or configuration files.

Physical backups have these characteristics:

- Consist of raw copies of database directories and folders.
- Output is more compact.
- Backups can include log and configuration files.
- Portable only to other machines with similar hardware and software.



Important

mysqldump requires at least the **Select** privilege for dumped tables, **SHOW VIEW** for dumped views, and **TRIGGER** for dumped triggers.

Useful Options

Option	Description
--add-drop-tables	Tells MariaDB to add a DROP TABLE statement before each CREATE TABLE statement.
--no-data	Dumps only the database structure, not the contents.
--lock-all-tables	No new record can be inserted anywhere in the database while the copy is finished. This option is very important to ensure backup integrity.
--add-drop-databases	Tells MariaDB to add a DROP DATABASE statement before each CREATE DATABASE statement.

Performing a physical backup

Several tools are available to perform physical backups, such as **ibbackup**, **cp**, **mysqlhotcopy**, and **lvm**.

A MariaDB physical backup task can use the known benefits of LVM snapshots. The following process will back up MariaDB using LVM.



Important

The key benefit of this methodology is that it is very quick, and keeps the downtime of the database short. This is a great argument for putting the database files on a dedicated LVM partition.

Verify where MariaDB files are stored:

```
[root@serverX ~]# mysqladmin variables | grep datadir
| datadir | /var/lib/mysql/ |
```

Verify which logical volume hosts this location:

```
[root@serverX ~]# df /var/lib/mysql
Filesystem      1K-blocks    Used Available Use% Mounted on
/dev/mapper/vg0-mariadb 51475068 7320316 41516928 15% /var/lib/mysql |
```

This shows that the volume group is **vg0** and the logical volume name is **mariadb**.

Verify how much space is available for the snapshot:

```
[root@serverX ~]# vgdisplay vg0 | grep Free
Free PE / Size      15321 / 61.29 GB |
```

This shows that 61.29 GB are available for a snapshot.

Connect to MariaDB, flush the tables to disk, and lock them (alternately, shut down the `mariadb` service):



Warning

This step is very important so that no new records are inserted into the database while the snapshot is created.

```
[root@serverX ~]# mysql -u root -p
MariaDB [(none)]> FLUSH TABLES WITH READ LOCK;
```



Warning

Do not close this session. As soon as the client disconnects, this lock is lifted. The database *must* remain locked until the LVM snapshot is created.

In another terminal session, create the LVM snapshot:

```
[root@serverX ~]# lvcreate -L20G -s -n mariadb-backup /dev/vg0/mariadb
```



Important

This snapshot needs to be large enough to hold the backup.

In the original MariaDB session, unlock the tables (or, bring the `mariadb` service up):

```
MariaDB [(none)]> UNLOCK TABLES;
```

The snapshot can now be mounted at an arbitrary location:

```
[root@serverX ~]# mkdir /mnt/snapshot
[root@serverX ~]# mount /dev/vg0/mariadb-backup /mnt/snapshot
```

From here, any standard file system backup can be used to store a copy of `/var/lib/mysql` as mounted under `/mnt/snapshot`.



Important

Do not forget to delete the snapshot once it has been backed up.

```
[root@serverX ~]# umount /mnt/snapshot
[root@serverX ~]# lvremove /dev/vg0/mariadb-backup
```

Restoring a backup

Logical restore

A logical restore can be done with the command `mysql`:

```
[root@serverX ~]# mysql -u root ❶ -p ❷ inventory ❸ < /backup/mariadb.dump ❹
```

- ❶ User to connect with to restore the MariaDB backup (generally root or some other superuser)
- ❷ Password for this user
- ❸ Selected database for restore backup
- ❹ Backup file

Physical restore

To do a physical restore, the mariadb service must be stopped:

```
[root@serverX ~]# systemctl stop mariadb
```

Verify where MariaDB files are stored:

```
[root@serverX ~]# mysqladmin variables | grep datadir  
| datadir | /var/lib/mysql/ |
```

Remove the actual content:

```
[root@serverX ~]# rm -rf /var/lib/mysql/*
```

From here, any standard file system restore can be used to restore a copy from backup to `/var/lib/mysql`.

R

References

`mysql`(1), `mysqladmin`(1), `mysqldump`(1), `lvcreate`(8), `lvremove`(1), and `vgdisplay`(8) man pages

- MariaDB Knowledge Base: *Backup and Restore Overview*
· <https://mariadb.com/kb/en/backup-and-restore-overview/>

Practice: Restoring a MariaDB Database from Backup

Guided exercise

In this lab, you will restore a database from a MariaDB logical backup.

Resources:	
Files:	/home/student/inventory.dump
Machines:	serverX

Outcomes:

A MariaDB database with restored data records.

Before you begin...

- Reset your serverX system.
- Log into and set up your serverX system.

```
[root@serverX ~]# lab mariadb-restore setup
```

You have been asked to restore from a logical backup to a blank MariaDB installation.

1. Connect to MariaDB as user root (blank password).

- 1.1.

```
[root@serverX ~]# mysql -u root
```

2. Create a database inventory.

- 2.1.

```
MariaDB [(none)]> create database inventory;
```

3. Restore from the logical backup.

- 3.1.

```
[root@serverX ~]# mysql -u root inventory < /home/student/inventory.dump
```

4. Verify restored data.

- 4.1. Connect again to MariaDB.

- 4.1.1.

```
[root@serverX ~]# mysql -u root
```

- 4.2. Connect to the inventory database.

- 4.2.1.

```
MariaDB [(none)]> use inventory;
```

- 4.3. Select all categories.

```
MariaDB [(inventory)]> SELECT * FROM category;
+----+-----+
| id | name  |
+----+-----+
|  1 | Networking |
|  2 | Servers  |
|  3 | Ssd      |
+----+-----+
3 rows in set (0.00 sec)
```

Lab: Configuring MariaDB Databases

Performance checklist

In this lab, you will install, configure, restore, and populate a MariaDB database.

Resources:	
Files:	/home/student/mariadb.dump
Machines:	desktopX and serverX

Outcomes:

A MariaDB service listening for remote connections. A backup file will be restored and three users with different privileges will be created.

Before you begin...

- Reset your **desktopX** system.
- Reset your **serverX** system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab mariadb setup
```

- Log into and set up your **desktopX** system.

```
[student@desktopX ~]$ lab mariadb setup
```

Your developers are working on a legacy system, and the MariaDB has stopped working. You have the logical backup file from last night. You are asked to install a new MariaDB database server and to restore from this backup.

A new ticket has been assigned to you to create new remote access accounts with the following information:

User	Password	Privileges
mary	mary_password	Select on all tables from legacy database
legacy	legacy_password	Select, insert, update, delete on all tables from legacy database
report	report_password	Select on all tables from legacy database

Since developers do not have access to the production environment, they asked you to insert new records in the **manufacturer** table:

Name	Seller	Phone number
HP	Joe Doe	+1 (432) 754-3509
Dell	Luke Skywalker	+1 (431) 219-4589

Name	Seller	Phone number
Lenovo	Darth Vader	+1 (327) 647-6784

1. Install MariaDB database groups.
2. Start and enable the MariaDB service.
3. Configure the firewall to accept remote connections.
4. Create the **legacy** database.
5. Restore the database backup (`/home/student/mariadb.dump`).
6. Create users according to the ticket.
7. Insert the new manufacturers.
8. From your desktopX system, validate your work.

8.1. `[student@desktopX ~]$ lab mariadb grade`

Solution

In this lab, you will install, configure, restore, and populate a MariaDB database.

Resources:	
Files:	/home/student/mariadb.dump
Machines:	desktopX and serverX

Outcomes:

A MariaDB service listening for remote connections. A backup file will be restored and three users with different privileges will be created.

Before you begin...

- Reset your **desktopX** system.
- Reset your **serverX** system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab mariadb setup
```

- Log into and set up your **desktopX** system.

```
[student@desktopX ~]$ lab mariadb setup
```

Your developers are working on a legacy system, and the MariaDB has stopped working. You have the logical backup file from last night. You are asked to install a new MariaDB database server and to restore from this backup.

A new ticket has been assigned to you to create new remote access accounts with the following information:

User	Password	Privileges
mary	mary_password	Select on all tables from legacy database
legacy	legacy_password	Select, insert, update, delete on all tables from legacy database
report	report_password	Select on all tables from legacy database

Since developers do not have access to the production environment, they asked you to insert new records in the **manufacturer** table:

Name	Seller	Phone number
HP	Joe Doe	+1 (432) 754-3509
Dell	Luke Skywalker	+1 (431) 219-4589
Lenovo	Darth Vader	+1 (327) 647-6784

1. Install MariaDB database groups.
 - 1.1. Install the MariaDB client on your desktopX machine.

```
[root@desktopX ~]# yum groupinstall mariadb-client -y
```

- 1.2. Install MariaDB database server on your serverX machine.

```
[root@serverX ~]# yum groupinstall mariadb -y
```

2. Start and enable the MariaDB service.

- 2.1. Start the MariaDB service.

```
[root@serverX ~]# systemctl start mariadb.service
```

- 2.2. Enable the MariaDB service.

```
[root@serverX ~]# systemctl enable mariadb.service
```

3. Configure the firewall to accept remote connections.

- 3.1.

```
[root@serverX ~]# firewall-cmd --permanent --add-service=mysql
```

- 3.2.

```
[root@serverX ~]# firewall-cmd --reload
```

4. Create the **legacy** database.

- 4.1. Connect to the MariaDB database server.

```
[root@serverX ~]# mysql -u root
```

- 4.2. Create the **legacy** database.

```
MariaDB [(none)]> create database legacy;  
MariaDB [(none)]> exit;
```

5. Restore the database backup (`/home/student/mariadb.dump`).

- 5.1.

```
[root@serverX ~]# mysql -u root legacy < /home/student/mariadb.dump
```

6. Create users according to the ticket.

- 6.1. Connect to the MariaDB database server.

```
[root@serverX ~]# mysql -u root
```

- 6.2. Create user mary.

```
MariaDB [(none)]> CREATE USER mary@'%' IDENTIFIED BY 'mary_password';
```

6.3. Create user legacy.

```
MariaDB [(none)]> CREATE USER legacy@'%' IDENTIFIED BY 'legacy_password';
```

6.4. Create user report.

```
MariaDB [(none)]> CREATE USER report@'%' IDENTIFIED BY 'report_password';
```

6.5. Grant privileges to mary.

```
MariaDB [(none)]> GRANT SELECT ON legacy.* TO mary@'%';
```

6.6. Grant privileges to legacy.

```
MariaDB [(none)]> GRANT SELECT,INSERT,UPDATE,DELETE ON legacy.* TO legacy@'%';
```

6.7. Grant privileges to report.

```
MariaDB [(none)]> GRANT SELECT ON legacy.* TO report@'%';
```

6.8. Flush privileges.

```
MariaDB [(none)]> FLUSH PRIVILEGES;
```

7. Insert the new manufacturers.

7.1. Select legacy database.

```
MariaDB [(none)]> USE legacy;
```

7.2. Describe table to verify attributes.

```
MariaDB [(legacy)]> DESCRIBE manufacturer;
```

Field	Type	Null	Key	Default	Extra
id	int(11)	NO	PRI	NULL	auto_increment
name	varchar(100)	NO		NULL	
seller	varchar(100)	YES		NULL	
phone_number	varchar(17)	YES		NULL	

```
4 rows in set (0.00 sec)
```

7.3. Insert HP as manufacturer.

```
MariaDB [(legacy)]> INSERT INTO manufacturer(name,seller,phone_number) values
('HP','Joe Doe','+1 (432) 754-3509');
```

7.4. Insert Dell as manufacturer.

```
MariaDB [(legacy)]> INSERT INTO manufacturer(name,seller,phone_number) values
('DELL','Luke Skywalker','+1 (431) 219-4589');
```

7.5. Insert Lenovo as manufacturer.

```
MariaDB [(legacy)]> INSERT INTO manufacturer(name,seller,phone_number) values
('Lenovo','Darth Vader','+1 (327) 647-6784');
```

8. From your desktopX system, validate your work.

8.1.

```
[student@desktopX ~]$ lab mariadb grade
```

Summary

Installing MariaDB

In this section, students learned how to:

- Describe relational databases.
- Describe database concepts.
- Perform a MariaDB installation.

Working with MariaDB Databases

In this section, students learned how to:

- Examine databases.
- Search in databases.
- Create databases.
- Use SQL.

Managing Database Users and Access Rights

In this section, students learned how to:

- Create users.
- Grant privileges.
- Revoke privileges.
- Flush privileges.

Creating and Restoring MariaDB Backups

In this section, students learned how to:

- Back up a MariaDB database.
- Restore a MariaDB database.



CHAPTER 10

PROVIDING APACHE HTTPD WEB SERVICE

Overview	
Goal	To configure Apache HTTPD to provide TLS-enabled websites and virtual hosts.
Objectives	<ul style="list-style-type: none">• Identify the key configuration files, log files, and content directories used by Apache <code>httpd</code>.• Configure Apache <code>httpd</code> to provide IP-based and name-based virtual hosts.• Configure Apache <code>httpd</code> to provide TLS-encrypted virtual hosts.• Configure Apache <code>httpd</code> to serve dynamic database-driven web content.
Sections	<ul style="list-style-type: none">• Configuring Apache HTTPD (and Practice)• Configuring and Troubleshooting Virtual Hosts (and Practice)• Configuring HTTPS (and Practice)• Integrating Dynamic Web Content (and Practice)
Lab	<ul style="list-style-type: none">• Configuring a Web Application

Configuring Apache HTTPD

Objectives

After completing this section, students should be able to identify the key configuration files, logfiles, and content directories used by Apache **httpd**.

Introduction to Apache HTTPD

Apache HTTPD is one of the most used web servers on the Internet. A web server is a daemon that speaks the **http(s)** protocol, a text-based protocol for sending and receiving objects over a network connection. The **http** protocol is sent over the wire in clear text, using port **80/TCP** by default (though other ports can be used). There is also a TLS/SSL encrypted version of the protocol called **https** that uses port **443/TCP** by default.

A basic http exchange has the client connecting to the server, and then requesting a resource using the **GET** command. Other commands like **HEAD** and **POST** exist, allowing clients to request just metadata for a resource, or send the server more information.

The following is an extract from a short http exchange:

```
GET /hello.html HTTP/1.1
Host: webapp0.example.com
User-Agent: Mozilla/5.0 (X11; Linux x86_64; rv:24.0) Gecko/20100101 Firefox/24.0
Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8
Accept-Language: en-US,en;q=0.5
Accept-Encoding: gzip, deflate
Connection: keep-alive
Cache-Control: max-age=0
```

The client starts by requesting a resource (the **GET** command), and then follows up with some extra headers, telling the server what types of encoding it can accept, what language it would prefer, etc. The request is ended with an empty line.

```
HTTP/1.1 200 OK
Date: Tue, 27 May 2014 09:57:40 GMT
Server: Apache/2.4.6 (Red Hat) OpenSSL/1.0.1e-fips mod_wsgi/3.4 Python/2.7.5
Content-Length: 12
Keep-Alive: timeout=5, max=82
Connection: Keep-Alive
Content-Type: text/plain; charset=UTF-8

Hello World!
```

The server then replies with a status code (**HTTP/1.1 200 OK**), followed by a list of headers. The **Content-Type** header is a mandatory one, telling the client what type of content is being sent. After the headers are done, the server sends an empty line, followed by the requested content. The length of this content must match the length indicated in the **Content-Length** header.

While the http protocol seems easy at first, implementing all of the protocol—along with security measures, support for clients not adhering fully to the standard, and support for dynamically

generated pages—is not an easy task. That is why most application developers do not write their own web servers, but instead write their applications to be run behind a web server like Apache HTTPD.

About Apache HTTPD

Apache HTTPD, sometimes just called “Apache” or **httpd**, implements a fully configurable and extendable web server with full **http** support. The functionality of **httpd** can be extended with *modules*, small pieces of code that plug into the main web server framework and extend its functionality.

On Red Hat Enterprise Linux 7 Apache HTTPD is provided in the *httpd* package. The *web-server* package group will install not only the *httpd* package itself, but also the *httpd-manual* package. Once *httpd-manual* is installed, and the **httpd.service** service is started, the *full* Apache HTTPD manual is available on <http://localhost/manual>. This manual has a complete reference of all the configuration directives for **httpd**, along with examples. This makes it an invaluable resource while configuring **httpd**.

Red Hat Enterprise Linux 7 also ships an environment group called *web-server-environment*. This environment group pulls in the *web-server* group by default, but has a number of other groups, like backup tool and database clients, marked as optional.

A default dependency of the *httpd* package is the *httpd-tools* package. This package includes tools to manipulate password maps and databases, tools to resolve IP addresses in logfiles to hostnames, and a tool (**ab**) to benchmark and stress-test web servers.

Basic Apache HTTPD configuration

After installing the *web-server* package group, or the *httpd* package, a default configuration is written to **/etc/httpd/conf/httpd.conf**.

This configuration serves out the contents of **/var/www/html** for requests coming in to any hostname over plain **http**.

The basic syntax of the **httpd.conf** is comprised of two parts: **Key Value** configuration directives, and HTML-like **<Blockname parameter>** blocks with other configuration directives embedded in them. Key/value pairs outside of a block affect the entire server configuration, while directives inside a block typically only apply to a part of the configuration indicated by the block, or when the requirement set by the block is met.


```

ServerRoot "/etc/httpd"1
Listen 802
Include conf.modules.d/*.conf3
User apache4
Group apache5
ServerAdmin root@localhost6
<Directory />7
    AllowOverride none
    Require all denied
</Directory>
DocumentRoot "/var/www/html"8
<Directory "/var/www">
    AllowOverride None
    Require all granted
</Directory>
<Directory "/var/www/html">
    Options Indexes FollowSymLinks
    AllowOverride None
    Require all granted
</Directory>
<IfModule dir_module>9
    DirectoryIndex index.html
</IfModule>
<Files ".ht*">10
    Require all denied
</Files>
ErrorLog "logs/error_log"11
LogLevel warn
<IfModule log_config_module>
    LogFormat "%h %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%{User-Agent}i\"" combined
    LogFormat "%h %l %u %t \"%r\" %>s %b" common
    <IfModule logio_module>
        LogFormat "%h %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%{User-Agent}i\" %I %0"
        combinedio
    </IfModule>
    CustomLog "logs/access_log" combined12
</IfModule>
AddDefaultCharset UTF-813
IncludeOptional conf.d/*.conf14

```

- ❶ This directive specifies where **httpd** will look for any files referenced in the configuration files with a relative path name.
- ❷ This directive tells **httpd** to start listening on port **80/TCP** on all interfaces. To only listen on select interfaces, the syntax "**Listen 1.2.3.4:80**" can be used for IPv4 or "**Listen [2001:db8::1]:80**" for IPv6.

Note: Multiple **listen** directives are allowed, but overlapping **listen** directives will result in a fatal error, preventing **httpd** from starting.

- ❸ This directive includes other files, as if they were inserted into the configuration file in place of the **Include** statement. When multiple files are specified, they will be sorted by filename in alphanumeric order before being included. Filenames can either be absolute, or relative to **ServerRoot**, and include wildcards such as *****.

- Note:** Specifying a nonexistent file will result in a fatal error, preventing **httpd** from starting.
- ④⑤ These two directives specify the user and group the **httpd** daemon should run as. **httpd** is always started as **root**, but once all actions that need **root** privileges have been performed—for example, binding to a port number under **1024**—privileges will be dropped and execution is continued as a nonprivileged user. This is a security measure.
 - ⑥ Some error pages generated by **httpd** can include a link where users can report a problem. Setting this directive to a valid email address will make the webmaster easier to contact for users. Leaving this setting at the default of **root@localhost** is not recommended.
 - ⑦ A **<Directory>** block sets configuration directives for the specified directory, and all descendent directories.

Common directives inside a **<Directory>** block include the following:

- **AllowOverride None:** **.htaccess** files will not be consulted for per-directory configuration settings. Setting this to any other setting will have a performance penalty, as well as the possible security ramifications.
- **Require All Denied:** **httpd** will refuse to serve content out of this directory, returning a **HTTP/1.1 403 Forbidden** error when requested by a client.
- **Require All Granted:** Allow access to this directory. Setting this on a directory outside of the normal content tree can have security implications.
- **Options [+|-]OPTIONS**...: Turn on (or off) certain options for a directory. For example, the **Indexes** option will show a directory listing if a directory is requested and no **index.html** file exists in that directory.
- ⑧ This setting determines where **httpd** will search for requested files. It is important that the directory specified here is both readable by **httpd** (both regular permissions and SELinux), and that a corresponding **<Directory>** block has been declared to allow access.
- ⑨ This block only applies its contents if the specified extension module is loaded. In this case, the **dir_module** is loaded, so the **DirectoryIndex** directive can be used to specify what file should be used when a directory is requested.
- ⑩ A **<Files>** block works just as a **<Directory>** block, but here options for individual (wildcarded) files is used. In this case, the block prevents **httpd** from serving out any security-sensitive files like **.htaccess** and **.htpasswd**.
- ⑪ This specifies to what file **httpd** should log any errors it encounters. Since this is a relative pathname, it will be prepended with the **ServerRoot** directive. In a default configuration, **/etc/httpd/logs** is a symbolic link to **/var/log/httpd/**.
- ⑫ The **CustomLog** directive takes two parameters: a file to log to, and a log format defined with the **LogFormat** directive. Using these directives, administrators can log exactly the information they need or want. Most log parsing tools will assume that the default **combined** format is used.
- ⑬ This setting adds a **charset** part to the **Content-Type** header for **text/plain** and **text/html** resources. This can be disabled with **AddDefaultCharset Off**.
- ⑭ This works the same as a regular include, but if no files are found, no error is generated.

Starting Apache HTTPD

httpd can be started from the **httpd.service** **systemd** unit.

```
[root@serverX ~]# systemctl enable httpd.service
[root@serverX ~]# systemctl start httpd.service
```

Once **httpd** is started, status information can be requested with **systemctl status -l httpd.service**. If **httpd** has failed to start for any reason, this output will typically give a clear indication of *why httpd* failed to start.

Network security

firewalld has two predefined services for **httpd**. The **http** service opens port **80/TCP**, and the **https** service opens port **443/TCP**.

```
[root@serverX ~]# firewall-cmd --permanent --add-service=http --add-service=https
[root@serverX ~]# firewall-cmd --reload
```

In a default configuration, **SELinux** only allows **httpd** to bind to a specific set of ports. This full list can be requested with the command **semanage port -l | grep '^http_'**. For a full overview of the allowed port contexts, and their intended usage, consult the **httpd_selinux(8)** man page from the *selinux-policy-devel* package.

Using an alternate document root

Content does not need to be served out of **/var/www/html**, but when changing the **DocumentRoot** setting, a number of other changes must be made as well:

- File system permissions: Any new **DocumentRoot** must be readable by the **apache** user or the **apache** group. In most cases, the **DocumentRoot** should *never* be writable by the **apache** user or group.
- SELinux: The default SELinux policy is restrictive as to what contexts can be read by **httpd**. The default context for web server content is **httpd_sys_content_t**. Rules are already in place to relabel **/srv/* /www/** with this context as well. To serve content from outside of these standard locations, a new context rule must be added with **semanage**.

```
[root@serverX ~]# semanage fcontext -a -t httpd_sys_content_t '/new/location(/.*)?'
```

Consult the **httpd_selinux(8)** man page from the *selinux-policy-devel* package for additional allowed file contexts and their intended usage.

Allowing write access to a DocumentRoot

In a default configuration only **root** has write access to the standard **DocumentRoot**. To allow web developers to write into the **DocumentRoot**, a number of approaches can be taken.

- Set a (default) ACL for the web developers on the **DocumentRoot**. For example, if all web developers are part of the **webmasters** group, and **/var/www/html** is used as the **DocumentRoot**, the following commands will give them write access:

```
[root@serverX ~]# setfacl -R -m g:webmasters:rwX /var/www/html
[root@serverX ~]# setfacl -R -m d:g:webmasters:rwX /var/www/html
```



Important

The uppercase **X** bit sets the executeable bit only on directories instead of directories and regular files. This is especially relevant when done in conjunction with a recursive action on a directory tree.

- Create the new **DocumentRoot** owned by the **webmasters** group, the sticky bit set.

```
[root@serverX ~]# mkdir -p -m 2775 /new/docroot  
[root@serverX ~]# chgrp webmasters /new/docroot
```

- A combination of the previous, with the other permissions closed off, and an ACL added for the **apache** group.



References

httpd(8) and **httpd_selinux(8)** man pages

httpd-manual package contents

Practice: Configuring a Web Server

Guided exercise

In this lab, you will configure a basic **httpd** web server to serve out a static page from the default location, as well as the Apache **httpd** manual.

Resources:	
Machines:	desktopX and serverX

Outcomes:

An Apache **httpd** web server running on your **serverX** machine, serving out a static page and the complete Apache **httpd** manual.

Before you begin...

- Reset your **desktopX** machine.
- Reset your **serverX** machine.

You have been asked to configure a basic web server on your **serverX** machine. This web server should serve out the text "Hello Class!" when the URL `http://serverX.example.com/` is requested.

To aid your end users in submitting bug reports, the default error page should include a **mailto:** reference to the email address **webmaster@serverX.example.com**.

Since your organization is planning on further customizing the behavior of this web server, the full Apache **httpd** manual should be available under `http://serverX.example.com/manual/`.

1. Begin by installing the *httpd* and *httpd-manual* packages.

1.1.

```
[student@serverX ~]$ sudo yum -y install httpd httpd-manual
```

2. Set the **ServerAdmin** directive for the main site configuration to point to **webmaster@serverX.example.com**.

- 2.1. Open `/etc/httpd/conf/httpd.conf` in a text editor with **root** privileges, and change the line that starts with **ServerAdmin** to the following:

```
ServerAdmin webmaster@serverX.example.com
```

3. Create the default content page.

- 3.1. Create the `/var/www/html/index.html` file with a text editor as user **root**, and add the following content:

```
Hello Class!
```

4. Start and enable the **httpd** service.

- 4.1.

```
[student@serverX ~]$ sudo systemctl start httpd.service
[student@serverX ~]$ sudo systemctl enable httpd.service
```
- 5. Open all the relevant ports for **http** on the firewall on **serverX**.
 - 5.1.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --add-service=http
[student@serverX ~]$ sudo firewall-cmd --reload
```
- 6. Test if you can access the new static page, as well as the Apache **httpd** manual, from your **desktopX** machine.
 - 6.1. From **desktopX**, point a web browser at `http://serverX.example.com`.
 - 6.2. From **desktopX**, point a web browser at `http://serverX.example.com/manual`.

Configuring and Troubleshooting Virtual Hosts

Objectives

After completing this section, students should be able to configure Apache **httpd** to provide IP-based and name-based virtual hosts.

Virtual hosts

Virtual hosts allow a single **httpd** server to serve content for multiple domains. Based on either the IP address of the server that was connected to, the hostname requested by the client in the http request, or a combination of both, **httpd** can use different configuration settings, including a different **DocumentRoot**.

Virtual hosts are typically used when it is not cost-effective to spin up multiple (virtual) machines to serve out many low-traffic sites; for example, in a shared hosting environment.

Configuring virtual hosts

Virtual hosts are configured using **<VirtualHost>** blocks inside the main configuration. To ease administration, these virtual host blocks are typically not defined inside `/etc/httpd/conf/httpd.conf`, but rather in separate `.conf` files in `/etc/httpd/conf.d/`.

The following is an example file, `/etc/httpd/conf.d/site1.conf`.

```
<Directory /srv/site1/www>❶
  Require all granted
  AllowOverride None
</Directory>

<VirtualHost 192.168.0.1:80>❷
  DocumentRoot /srv/site1/www❸
  ServerName site1.example.com❹
  ServerAdmin webmaster@site1.example.com❺
  ErrorLog "logs/site1_error_log"❻
  CustomLog "logs/site1_access_log" combined❼
</VirtualHost>
```

- ❶ This block provides access to the **DocumentRoot** defined further down.
- ❷ This is the main tag of the block. The **192.168.0.1:80** part indicates to **httpd** that this block should be *considered* for all connections coming in on that IP/port combination.
- ❸ Here the **DocumentRoot** is being set, but only for within this virtual host.
- ❹ This setting is used to configure *name-based* virtual hosting. If multiple **<VirtualHost>** blocks are declared for the same IP/port combination, the block that matches **ServerName** with the **hostname**: header sent in the client **http** request will be used.

There can be exactly zero or one **ServerName** directives inside a single **<VirtualHost>** block. If a single virtual host needs to be used for more than one domain name, one or more **ServerAlias** statements can be used.

- 5 To help with sorting mail messages regarding the different websites, it is helpful to set unique **ServerAdmin** addresses for all virtual hosts.
- 6 The location for all error messages related to this virtual host.
- 7 The location for all access messages regarding this virtual host.

If a setting is not made explicitly for a virtual host, the same setting from the main configuration will be used.

Name-based vs. IP-based virtual hosting

By default, every virtual host is an IP-based virtual host, sorting traffic to the virtual hosts based on what IP address the client had connected to. If there are multiple virtual hosts declared for a single IP/port combination, the **ServerName** and **ServerAlias** directives will be consulted, effectively enabling name-based virtual hosting.

Wildcards and priority

The IP address part of a **<VirtualHost>** directive can be replaced with one of two wildcards: **_default_** and *****. Both have exactly the same meaning: "Match Anything".

When a request comes in, **httpd** will first try to match against virtual hosts that have an explicit IP address set. If those matches fail, virtual hosts with a wildcard IP address are inspected. If there is still no match, the "main" server configuration is used.



Important

A **<VirtualHost *:80>** will always match for regular http traffic on port **80/TCP**, effectively disabling the main server configuration from ever being used for traffic on port **80/TCP**.

If no exact match has been found for a **ServerName** or **ServerAlias** directive, and there are multiple virtual hosts defined for the IP/port combination the request came in on, the *first* virtual host that matches an IP/port is used, with *first* being seen as the order in which virtual hosts are defined in the configuration file.

When using multiple ***.conf** files, they will be included in alphanumeric sorting order. To create a catch-all (default) virtual host, the configuration file should be named something like **00-default.conf** to make sure that it is included before any others.

Troubleshooting virtual hosts

When troubleshooting virtual hosts, there are a number of approaches that can help.

- Configure a separate **DocumentRoot** for each virtual host, with identifying content.
- Configure separate logfiles, both for error logging and access logging, for each virtual host.
- Evaluate the order in which the virtual host definitions are parsed by **httpd**. Included files are read in alphanumeric sort order based on their filenames.
- Disable virtual hosts one by one to isolate the problem. Virtual host definitions can be commented out of the configuration file(s), and include files can be temporarily renamed to something that does not end in **.conf**.
- **journalctl UNIT=httpd.service** can isolate log messages from just the **httpd.service** service.



References

httpd(8) man page

httpd-manual package contents

Practice: Configuring a Virtual Host

Guided exercise

In this lab, you will configure a new web server to serve out content for multiple virtual hosts.

Resources:

Machines:	desktopX and serverX
------------------	-----------------------------

Outcomes:

A new web server running on **serverX**, serving out content for `wwwX.example.com` from `/srv/www0.example.com/www/`, and all other domains from `/srv/default/www/`.

Before you begin...

- Reset your **serverX** machine.
- Reset your **desktopX** machine.

Over the past few years, your company has been spinning up many web servers for new projects. Unfortunately, there was no structure or coordination between the various projects.

In an effort to clean up the mess, you have been asked to consolidate these various web servers into one, serving out the different domains using name-based virtual hosting.

For now, you will only have to set up a *default* virtual host that serves out a placeholder site from `/srv/default/www/`, and a virtual host for `wwwX.example.com` that serves out content from `/srv/wwwX.example.com/www`.

DNS CNAME records for the relevant domains have already been converted to point at your **serverX** machine.

- 1. Start by installing the *httpd* package.

- 1.1.

```
[student@serverX ~]$ sudo yum install httpd
```

- 2. Create the content directories. The placeholder site should have an **index.html** file that reads:

```
Coming Soon!
```

The `wwwX.example.com` site should have an **index.html** that reads:

```
wwwX
```

- 2.1. Create the directories.

```
[student@serverX ~]$ sudo mkdir -p /srv/{default,wwwX.example.com}/www
```

- 2.2. Create the `index.html` files using a text editor. `/srv/default/www/index.html` gets the "Coming Soon!" text, and the `/srv/wwwX.example.com/www/index.html` file should read "wwwX".
- 2.3. The SELinux policy is already configured for the correct file contexts for any files under `srv/*www`, but you will still have to reset the context on the files you just created.

```
[student@serverX ~]$ sudo restorecon -Rv /srv/
```

- 3. Create a new virtual host definition for the `_default_:80` virtual host. This virtual host should serve out content from `/srv/default/www/`, and log to `logs/default-vhost.log` using the combined format.
 - 3.1. Create a new file called `/etc/httpd/conf.d/00-default-vhost.conf`. Give it the following content:

```
<VirtualHost _default_:80>
  DocumentRoot /srv/default/www
  CustomLog "logs/default-vhost.log" combined
</VirtualHost>
```

- 3.2. Since in a default configuration, `httpd` blocks access to all directories, you will need to open up the content directory for your default vhost. Add the following block to `/etc/httpd/conf.d/00-default-vhost.conf`.

```
<Directory /srv/default/www>
  Require all granted
</Directory>
```

- 4. Create a new virtual host definition for a `wwwX.example.com` virtual host in `/etc/httpd/conf.d/01-wwwX.example.com-vhost.conf`. This virtual host should respond to requests for both `wwwX.example.com` and `wwwX`, serve out content from `/srv/wwwX.example.com/www`, and store logs in `logs/wwwX.example.com.log`.
 - 4.1. Create the file `/etc/httpd/conf.d/01-wwwX.example.com-vhost.conf` with the following contents:

```
<VirtualHost *:80>
  ServerName wwwX.example.com
  ServerAlias wwwX
  DocumentRoot /srv/wwwX.example.com/www
  CustomLog "logs/wwwX.example.com.log" combined
</VirtualHost>

<Directory /srv/wwwX.example.com/www>
  Require all granted
</Directory>
```

- 5. Start and enable the `httpd` service.

□ 5.1.

```
[student@serverX ~]$ sudo systemctl start httpd.service
[student@serverX ~]$ sudo systemctl enable httpd.service
```

□ 6. Open up the firewall on **serverX** to allow traffic to the **httpd** service.

□ 6.1.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --add-service=http
[student@serverX ~]$ sudo firewall-cmd --reload
```

□ 7. From your **desktopX** system, use a web browser to visit the following URLs; the first two should respond with the "**wwwX**" text, while the last two should respond with "**Coming Soon!**".

- `http://wwwX.example.com`
- `http://wwwX`
- `http://serverX.example.com`
- `http://172.25.X.11`

Configuring HTTPS

Objectives

After completing this section, students should be able to configure Apache `httpd` to provide TLS-encrypted virtual hosts.

Transport Layer Security (TLS)

Transport Layer Security (TLS) is a method for encrypting network communications. TLS is the successor to *Secure Sockets Layer* (SSL). TLS allows a client to verify the identity of the server and, optionally, allows the server to verify the identity of the client.

TLS is based around the concepts of *certificates*. A certificate has multiple parts: a public key, server identity, and a signature from a *certificate authority*. The corresponding private key is never made public. Any data encrypted with the private key can only be decrypted with the public key, and vice versa.

During the initial *handshake*, when setting up the encrypted connection, the client and server agree on a set of encryption ciphers supported by both the server and the client, and they exchange bits of random data. The client uses this random data to generate a *session key*, a key that will be used for much faster *symmetric* encryption, where the same key is used for both encryption and decryption. To make sure that this key is not compromised, it is sent to the server encrypted with the server's public key (part of the server certificate).

The following diagram shows a (simplified) version of a TLS handshake.

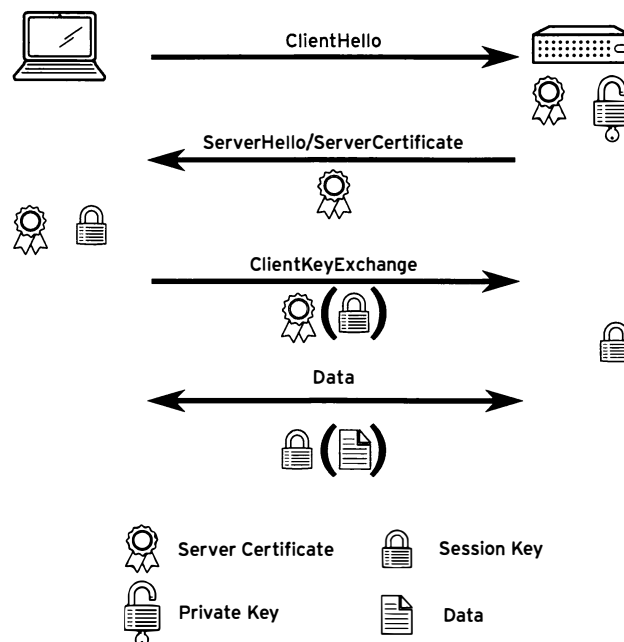


Figure 10.1: A simplified representation of a TLS handshake

1. The client initiates a connection to the server with a **ClientHello** message. As part of this message, the client sends a 32-byte random number including a timestamp, and a list of encryption protocols and ciphers supported by the client.
2. The server responds with a **ServerHello** message, containing another 32-byte random number with a timestamp, and the encryption protocol and ciphers the client should use.

The server also sends the server certificate, which consists of a public key, general server identity information like the FQDN, and a signature from a trusted *certificate authority* (CA). This certificate *can* also include the public certificates for all certificate authorities that have signed the certificate, up to a root CA.

3. The client verifies the server certificate by checking if the supplied identity information matches, and by verifying all signatures, checking if they are made by a CA trusted by the client.

If the certificate verifies, the client creates a *session key* using the random numbers previously exchanged. The client then encrypts this session key using the public key from the server certificate, and sends it to the server using a **ClientKeyExchange** message.

4. The server decrypts the session key, and the client and server both start encrypting and decrypting all data send over the connection using the session key.



Note

This is a simplification of the actual protocol; for example, the actual session key never gets transmitted with a lot cipher suites, not even in encrypted form. The server and client both create a *pre-master key* which gets exchanged, and both the server and client calculate the actual session key from that one.

During the negotiations, both the server and client also use a variety of methods to ensure against replay and man-in-the-middle attacks.

Configuring TLS certificates

To configure a virtual host with TLS, multiple steps must be completed:

1. Obtain a (signed) certificate.
2. Install Apache HTTPD extension modules to support TLS.
3. Configure a virtual host to use TLS, using the certificates obtained earlier.

Obtaining a certificate

When obtaining a certificate, there are two options: creating a self-signed certificate (a certificate signed by itself, not an actual CA), or creating a *certificate request* and having a reputable CA sign that request so it becomes a certificate.

The *crypto-utils* package contains a utility called **genkey** that supports both methods. To create a certificate (signing request) with **genkey**, run the following command, where **<FQDN>** is the fully qualified domain name clients will use to connect to your server:

```
[root@serverX ~]# genkey <FQDN>
```

During the creation, **genkey** will ask for the desired key size (choose at least **2048** bits), if a signing request should be made (answering no will create a self-signed certificate), whether the private key should be protected with a passphrase, and general information about the identity of the server.

After the process has completed, a number of files will be generated:

- **/etc/pki/tls/private/<fqdn>.key**: This is the private key. The private key should be kept at **0600** or **0400** permissions, and an SELinux context of **cert_t**. This key file should never be shared with the outside world.
- **/etc/pki/tls/certs/<fqdn>.0.csr**: This file is only generated if you requested a signing request. This is the file that you send to your CA to get it signed. You *never* need to send the private key to your CA.
- **/etc/pki/tls/certs/<fqdn>.crt**: This is the public certificate. This file is only generated when a self-signed certificate is requested. If a signing request was requested and sent to a CA, this is the file that will be returned from the CA. Permissions should be kept at **0644**, with an SELinux context of **cert_t**.

Install Apache HTTPD modules

Apache HTTPD needs an extension module to be installed to activate TLS support. On Red Hat Enterprise Linux 7, you can install this module using the *mod_ssl* package.

This package will automatically enable **httpd** for a default virtual host listening on port **443/TCP**. This default virtual host is configured in the file **/etc/httpd/conf.d/ssl.conf**.

Configure a virtual host with TLS

Virtual hosts with TLS are configured in the same way as regular virtual hosts, with some additional parameters. It is possible to use name-based virtual hosting with TLS, but some older browsers are not compatible with this approach.

The following is a simplified version of **/etc/httpd/conf.d/ssl.conf**:

```

Listen 443 https1
SSLPassPhraseDialog exec:/usr/libexec/httpd-ssl-pass-dialog2
SSLSessionCache shmcb:/run/httpd/sslcache(512000)
SSLSessionCacheTimeout 300
SSLRandomSeed startup file:/dev/urandom 256
SSLRandomSeed connect builtin
SSLCryptoDevice builtin

<VirtualHost _default_:443>3
    ErrorLog logs/ssl_error_log
    TransferLog logs/ssl_access_log
    LogLevel warn

    SSLEngine on4
    SSLProtocol all -SSLv25
    SSLCipherSuite HIGH:MEDIUM:!aNULL:!MD56
    SSLCertificateFile /etc/pki/tls/certs/localhost.crt7
    SSLCertificateKeyFile /etc/pki/tls/private/localhost.key8
    CustomLog logs/ssl_request_log \
        "%t %h %{SSL_PROTOCOL}x %{SSL_CIPHER}x \"%r\" %b"
</VirtualHost>

```

- ❶ This directive instructs **https** to listen on port **443/TCP**. The second argument (**https**) is optional, since **https** is the default protocol for port **443/TCP**.
- ❷ If the private key is encrypted with a passphrase, **httpd** needs a method of requesting a passphrase from a user at the console at startup. This directive specifies what program to execute to retrieve that passphrase.
- ❸ This is the virtual host definition for a catch-all virtual host on port **443/TCP**.
- ❹ This is the directive that actually turns on **TLS** for this virtual host.
- ❺ This directive specifies the list of protocols that **httpd** is willing to speak with clients. For added security, the older, unsafe **SSLv3** protocol should also be disabled:

```
SSLProtocol all -SSLv2 -SSLv3
```

- ❻ This directive lists what encryption ciphers **httpd** is willing to use when communicating with clients. The selection of ciphers can have big impacts on both performance and security.
- ❼ This directive instructs **httpd** where it can read the certificate for this virtual host.
- ❽ This directive instructs **httpd** where it can read the private key for this virtual host. **httpd** reads all private keys before privileges are dropped, so file permissions on the private key can remain locked down.

If a certificate signed by an CA is used, and the certificate itself does not have copies of all the CA certificates used in signing, up to a root CA, embedded in it, the server will also need to provide a *certificate chain*, a copy of all CA certificates used in the signing process concatenated together. The **SSLCertificateChainFile** directive is used to identify such a file.

When defining a new TLS-encrypted virtual host, it is not needed to copy the entire contents of **ssl.conf**. Only a **<VirtualHost>** block with the **SSLEngine On** directive, and configuration for certificates, is strictly needed. The following is an example of a name-based TLS virtual host:

```
<VirtualHost *:443>
  ServerName demo.example.com
  SSLEngine on
  SSLCertificateFile /etc/pki/tls/certs/demo.example.com.crt
  SSLCertificateKeyFile /etc/pki/tls/private/demo.example.com.key
  SSLCertificateChainFile /etc/pki/tls/certs/example-ca.crt
</VirtualHost>
```

This example misses some important directives such as **DocumentRoot**; these will be inherited from the main configuration.



Warning

Not defining what protocols and ciphers can be used will result in **httpd** using default options for these. **httpd** defaults are not considered secure, and it is highly recommended to restrict both to a more secure subset.

Configuring forward secrecy

If a weaker encryption cipher has been used, and the private key of the server has been compromised—for example, after a server break-in or a bug in the crypto code—an attacker could possibly decrypt a recorded session.

Protecting against these types of attacks is called ensuring *forward secrecy*. Forward secrecy can be established by carefully tuning the allowed ciphers in the **SSLCipherSuite** directive, and having the server always select the most preferred cipher from this list that both the server and the client support.

The following is an example that, at the date of publication, was considered the best set of ciphers to allow. This list prioritizes ciphers that perform the initial session key exchange using elliptic curve Diffie-Hellman (EECDH) algorithms. Using Diffie-Hellman, the actual session key is never transmitted, but rather calculated by both sides.

```
SSLCipherSuite "EECDH+ECDSA+AESGCM EECDH+aRSA+AESGCM EECDH+ECDSA+SHA384 EECDH+ECDSA
+SHA256 EECDH+aRSA+SHA384 EECDH+aRSA+SHA256 EECDH+aRSA+RC4 EECDH EDH+aRSA RC4 !aNULL !
eNULL !LOW !3DES !MD5 !EXP !PSK !SRP !DSS"
SSLHonorCipherOrder On
```

The **SSLHonorCipherOrder On** directive instructs **httpd** to always prefer ciphers listed earlier on in the **SSLCipherSuite** list, regardless of the client preference.



Important

Security research is an always ongoing arms race. It is recommended that administrators re-evaluate their selected ciphers on a regular basis.

Configuring HTTP Strict Transport Security (HSTS)

A common misconfiguration, and one that will result in warnings in most modern browsers, is having a web page that is served out over **https** include resources served out over clear-text **http**.

To protect against this type of misconfiguration, add the following line inside a **<VirtualHost>** block that has TLS enabled:

```
Header always set Strict-Transport-Security "max-age=15768000"
```

Sending this extra header informs clients that they are not allowed to fetch any resources for this page that are not served using TLS.

Another possible issue comes from clients connecting over **http** to a resource they should have been using **https** for.

Simply not serving any content over **http** would alleviate this issue, but a more subtle approach is to automatically redirect clients connecting over **http** to the same resource using **https**.

To set up these redirects, configure a **http** virtual host for the same **ServerName** and **ServerAlias** as the TLS protected virtual host (a catch-all virtual host can be used), and add the following lines inside the **<VirtualHost *:80>** block:

```
RewriteEngine on
RewriteRule ^(/.*)$ https://%{HTTP_POST}$1 [redirect=301]
```

The **RewriteEngine on** directive turns on the URL rewrite module for this virtual host, and the **RewriteRule** matches any resource (**^(/.*)\$**) and redirects it using a **http Moved**

Permanently message (**[redirect=301]**) to the same resource served out over https. The **%{HTTP_HOST}** variable uses the hostname that was requested by the client, while the **\$1** part is a back-reference to whatever was matched between the first set of parentheses in the regular expression.



References

httpd(8) man page

httpd-manual contents

Qualys SSL Labs: SSL/TLS Deployment Best Practices
<https://www.ssllabs.com/projects/best-practices/>

Practice: Configuring a TLS-enabled Virtual Host

Guided exercise

In this lab, you will configure TLS-encrypted virtual hosts.

Resources:	
Files:	<ul style="list-style-type: none"> • http://classroom.example.com/pub/example-ca.crt • http://classroom.example.com/pub/tls/certs/wwwX.crt • http://classroom.example.com/pub/tls/private/wwwX.key • http://classroom.example.com/pub/tls/certs/webappX.crt • http://classroom.example.com/pub/tls/private/webappX.key
Machines:	desktopX and serverX

Outcomes:

A web server configured with two virtual hosts, **wwwX.example.com** and **webappX.example.com**, both protected by TLS.

Before you begin...

- Reset your **serverX** system.
- Reset your **desktopX** system.

Your company has decided to start selling Jim Whitehurst action figures online. Since most Red Hat fans enjoy privacy and security, the website will need to be protected with TLS.

You have been asked to configure a web server on your **serverX** machine to host this site. This web server will need to host two virtual hosts: <https://wwwX.example.com> and <https://webappX.example.com>. The nonencrypted version of these two sites should send browsers an automatic redirect to the encrypted version.

Certificates and private keys for these two sites have already been provided. The certificates can be downloaded from <http://classroom.example.com/pub/tls/certs/{www,webapp}X.crt>, and the private keys can be found at <http://classroom.example.com/pub/tls/private/{www,webapp}X.key>. The public part of the signing CA can be found at <http://classroom.example.com/pub/example-ca.crt>.

Content for these sites should be served out of `/srv/wwwX/www` and `/srv/webappX/www`, respectively. Since your web designers are currently on a two-week lunch break, you will have to provide temporary content that uniquely identifies each host yourself.

Custom logfiles are not required for now.

- ❑ 1. Install both the `httpd` and `mod_ssl` packages.

□ 1.1.

```
[student@serverX ~]$ sudo yum install httpd mod_ssl
```

□ 2. Create the content directories, with identifying content and appropriate SELinux contexts.

□ 2.1. Create the two directories.

```
[student@serverX ~]$ sudo mkdir -p /srv/{www,webapp}X/www
```

□ 2.2. In both content directories, create an **index.html** file with distinct content.

```
[student@serverX ~]$ sudo vim /srv/wwwX/www/index.html
[student@serverX ~]$ sudo vim /srv/webappX/www/index.html
```

□ 2.3. Reset the SELinux context on your new directories.

```
[student@serverX ~]$ sudo restorecon -Rv /srv/
```

□ 3. Download all the needed certificates and private keys to their correct locations with their correct permissions.

□ 3.1. Download the CA certificate used to sign your certificates.

```
[student@serverX ~]$ cd /etc/pki/tls/certs
[student@serverX certs]$ sudo wget http://classroom.example.com/pub/
example-ca.crt
```

□ 3.2. While still in the **certs** directory, download the two certificates for your virtual hosts.

```
[student@serverX certs]$ sudo wget http://classroom.example.com/pub/tls/
certs/wwwX.crt
[student@serverX certs]$ sudo wget http://classroom.example.com/pub/tls/
certs/webappX.crt
```

□ 3.3. Switch to the **private** directory and download the private keys. Do not forget to set the permissions on the private keys to **0600**.

```
[student@serverX certs]$ cd /etc/pki/tls/private
[student@serverX private]$ sudo wget http://classroom.example.com/pub/tls/
private/wwwX.key
[student@serverX private]$ sudo wget http://classroom.example.com/pub/tls/
private/webappX.key
[student@serverX private]$ sudo chmod 0600 w*X.key
```

□ 4. Configure the TLS name-based virtual host for your **wwwX.example.com** domain in a new file called **/etc/httpd/conf.d/wwwX.conf**. You can use the existing **/etc/httpd/conf.d/ssl.conf** as a template, but if you do, do not forget to strip out all the content outside of the **<VirtualHost>** block.

Do not forget to add an automatic redirect from the non-TLS-based http site to the TLS-encrypted https site.

- 4.1. Create `/etc/httpd/conf.d/wwwX.conf` with the following content:

```
<VirtualHost *:443>
  ServerName wwwX.example.com
  SSLEngine On
  SSLProtocol all -SSLv2 -SSLv3
  SSLCipherSuite HIGH:MEDIUM:!aNULL:!MD5
  SSLHonorCipherOrder on
  SSLCertificateFile /etc/pki/tls/certs/wwwX.crt
  SSLCertificateKeyFile /etc/pki/tls/private/wwwX.key
  SSLCertificateChainFile /etc/pki/tls/certs/example-ca.crt
  DocumentRoot /srv/wwwX/www
</VirtualHost>
```

- 4.2. Add a `<Directory>` block for `/srv/wwwX/www` to `/etc/httpd/conf.d/wwwX.conf` like the following:

```
<Directory /srv/wwwX/www>
  Require all granted
</Directory>
```

- 4.3. To accomplish the automatic redirect from http to https, add the following block to `/etc/httpd/conf.d/wwwX.conf`:

```
<VirtualHost *:80>
  ServerName wwwX.example.com
  RewriteEngine on
  RewriteRule ^(/.*)$ https://%{HTTP_HOST}$1 [redirect=301]
</VirtualHost>
```

- 5. Configure the `webappX.example.com` virtual host by copying the configuration for your `wwwX.example.com` virtual host, and changing every occurrence of `wwwX` to `webappX`.

- 5.1. Copy the configuration to `/etc/httpd/conf.d/webappX.conf`.

```
[student@serverX ~]$ sudo cp /etc/httpd/conf.d/{www,webapp}X.conf
```

- 5.2. Replace every occurrence of `wwwX` with `webappX` in the new configuration file.

```
[student@serverX ~]$ sudo sed -i 's/wwwX/webappX/g' /etc/httpd/conf.d/
webappX.conf
```

- 6. Start and enable the `httpd.service`, and open the relevant firewall ports.

- 6.1. Start and enable `httpd.service`.

```
[student@serverX ~]$ sudo systemctl start httpd.service
[student@serverX ~]$ sudo systemctl enable httpd.service
```

- 6.2. Open both the **http** and **https** ports on the firewall.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --add-service=http --
add-service=https
[student@serverX ~]$ sudo firewall-cmd --reload
```

- 7. Test your new configuration from your **desktopX** system. You will have to import the `http://classroom.example.com/pub/example-ca.crt` into the list of trusted CA certificates for your browser as part of this process.

Perform all of the following steps on your **desktopX** system.

- 7.1. Download the **example-ca.crt** certificate to your home directory.

```
[student@desktopX ~]$ wget http://classroom.example.com/pub/example-ca.crt
```

- 7.2. Launch **firefox** and open the **Edit > Preferences** dialog. Navigate to the **Advanced > Certificates** tab.
- 7.3. Click **View Certificates**, then use the **Import** button. Navigate to the file you just downloaded and click **Open**. In the resulting dialog, check **Trust this CA to identify websites** and click **OK**.

Close all open dialogs.

- 7.4. Point your browser at both `http://wwwX.example.com` and `http://webappX.example.com`. Both should redirect to their **https** counterpart automatically, without a certificate warning.



Note

When troubleshooting a web server using **firefox**, it can be useful to empty the cache from within the **Preferences** dialog. The **Clear Now** button can be found in the **Advanced > Network** tab. If the cache is not cleared in between server restarts, **firefox** might show old, outdated information.

- 8. *Bonus question:*

Without further configuration, a visit to `http://serverX.example.com` will also result in a redirect to **https**. Why is this, and how could you prevent this from happening?

Answer:

This happens because there is an explicit catch-all virtual host defined for ***:80**, resulting in the first virtual host for ***:80** being used as a default virtual host. Since this is the virtual host for your **webappX** domain, the redirect rule is included.

This can be solved by defining either a `<VirtualHost _default_:80>` block, or by defining a `<VirtualHost *:80>` block in a location where

it will be parsed before any other virtual hosts; for example, before the includes in `/etc/httpd/conf/httpd.conf` or as a separate file in `/etc/httpd/conf.d/00-default.conf`.

Integrating Dynamic Web Content

Objectives

After completing this section, students should be able to configure Apache **httpd** to serve dynamic database-driven web content.

Dynamic content

Most modern websites do not consist of purely static content. Most content served out is actually generated dynamically, on demand. Integrating dynamic content with Apache HTTPD can be done in numerous ways. This section describes a number of the most common ways, but more ways exist.

Common Gateway Interface

One of the oldest forms of generating dynamic content is by using *Common Gateway Interface* (CGI). When a CGI resource is requested, **httpd** does not simply read the resource and serve it out; instead, it executes the resource as a process, and serves the **stdout** of that process. Although CGI resources are mostly written in scripting languages like Perl, it is also quite common for CGI resources to be compiled C programs, or Java executables.

Information from the request (including client information) is made available to the CGI program using environment variables.

Configuring httpd for CGI

To have **httpd** treat a location as CGI executables, the following syntax is used in the **httpd** configuration.

```
ScriptAlias /cgi-bin/ "/var/www/cgi-bin/"
```

This instructs **httpd** to redirect any request for files under the **/cgi-bin/** URI to the **/var/www/cgi-bin/** directory, and treat the files in that directory as executable scripts.

A number of caveats exist when using CGI:

- CGI scripts will be executed as the **apache** user and group.
- CGI scripts should be executable by the **apache** user and group.
- CGI scripts should have the **httpd_sys_script_exec_t** SELinux context.
- The CGI directory should have **Options None**, and access should be granted using a normal **<Directory>** block.

Serving dynamic PHP content

A popular method of providing dynamic content is using the PHP scripting language. While PHP scripts can be served using old-fashioned CGI, both performance and security can be improved by having **httpd** run a PHP interpreter internally.

By installing the *php* package, a special **mod_php** module is added to **httpd**. The default configuration for this module adds the following lines to the main **httpd** configuration:


```
<FilesMatch \.php$>
  SetHandler application/x-httpd-php
</FilesMatch>
DirectoryIndex index.php
```

The **<FilesMatch>** block instructs **httpd** to use **mod_php** for any file with a name ending in **.php**, and the **DirectoryIndex** directive adds **index.php** to the list of files that will be sought when a directory is requested.

Serving dynamic Python content

Also popular is generating dynamic content using Python scripts. Python scripts can be served out using regular CGI, but both **python** and **httpd** also support a newer protocol: *Web Server Gateway Interface* (WSGI).

WSGI support can be added to **httpd** by installing the *mod_wsgi* package.

Unlike the **mod_php** or CGI approach, WSGI does not start a new script/interpreter for every request. Instead, a main application is started, and all requests are routed into that application.

Configuring **httpd** to support a WSGI application takes two steps:

1. Install the *mod_wsgi* package.
2. Add a **WSGIScriptAlias** line to a virtual host definition.

The following is an example of a **WSGIScriptAlias** directive, which sends all requests for `http://servername/myapp` and any resources below it to the WSGI application `/srv/myapp/www/myapp.py`:

```
WSGIScriptAlias /myapp/ /srv/myapp/www/myapp.py
```

WSGI applications should be executable by the **apache** user and group, and their SELinux contexts should be set to **httpd_sys_content_t**.

Database connectivity

Most web applications will need to store and retrieve persistent data. A common approach to this is to store the data in a database such as MariaDB or PostgreSQL.

When the database is running on the same host as the web server, and the database is using a standard network port, SELinux will allow the network connection from the web application to happen.

When a database on a remote host is used, the SELinux Boolean **httpd_can_network_connect_db** must be set to **1** to allow the connection.

When a network connection to another needs to be made from within the web application, and the target is not a well-known database port, the SELinux Boolean **httpd_can_network_connect** must be set to **1**.

Various other SELinux Booleans can also affect the way in which web applications are executed by **httpd**.



References

httpd(8) and **httpd_php_selinux(8)** man pages

httpd-manual package contents

/usr/share/doc/mod_wsgi-*/README

Practice: Configuring a Web Application

Guided exercise

In this lab, you will configure your **serverX** system to serve a PHP application that uses a MariaDB database backend.

Resources:	
Files:	<code>/var/www/html/index.html</code>
Machines:	desktopX and serverX

Outcomes:

A working PHP application with a database backend running on **serverX**.

Before you begin...

- Reset your **desktopX** system.
- Reset your **serverX** system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab phpdb setup
```

Your company is trying to push into the collectible card game market. One of the major focuses in that push is to provide an online database of all cards in the popular game “Tragic: The Saddening”.

Your web developers are hard at work, and have provided you with an early alpha version of the web interface in `/var/www/html/index.php`.

Your database administrators, not wanting to be outdone by the web developers, have populated a **MariaDB** database, ready for use.

One of your former coworkers was tasked with configuring the web server for this project. Sadly, after a number of weeks trying to get the code to run, that coworker was let go.

It's now up to you to get this project on the road again. Both the database server and the web server are already running on **serverX**, but clients have trouble connecting to `http://serverX.example.com`.

Your web developers have delivered bug-free code as usual, and the DBA has delivered a working database, with correct records.

1. Replicate the issue of users not being able to connect to `serverX.example.com`.
 - 1.1. On your **desktopX** machine, start an instance of **firefox** and point it at `serverX.example.com`.
2. That “Unable to connect” message has two likely causes: **httpd.service** not running on **serverX**, or a firewall issue. Investigate both.
 - 2.1. Check if **httpd.service** is running on **serverX**.

```
[student@serverX ~]$ sudo systemctl status -l httpd.service
```

- 2.2. Investigate if the **http** service is opened for the default zone on **serverX**.

```
[student@serverX ~]$ sudo firewall-cmd --list-all
```

- 3. Configure the firewall on **serverX** to allow **http** traffic.

- 3.1.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --add-service=http  
[student@serverX ~]$ sudo firewall-cmd --reload
```

- 4. Test if opening up the firewall solved your issue.

- 4.1. From **desktopX**, point a browser at `http://serverX.example.com`.

- 5. That test page is typically shown when **httpd** cannot find an index page. Verify that there is indeed a file named `/var/www/html/index.php`, and that it is readable by the **httpd** process.

- 5.1.

```
[student@serverX ~]$ ls -lZ /var/www/html/index.php
```

- 6. **httpd** only searches for **.php** files when *php* is properly installed. Verify that this is the case.

- 6.1.

```
[student@serverX ~]$ yum list php
```

- 6.2. Install the *php* package.

```
[student@serverX ~]$ sudo yum install php
```

- 6.3. Restart the **httpd.service** on **serverX** to activate the new PHP modules.

```
[student@serverX ~]$ sudo systemctl restart httpd.service
```

- 7. From your **desktopX** system, use **firefox** to check if `http://serverX.example.com` is now working.

- 8. That empty page does not look good. Investigate the issue.

- 8.1. Check the **httpd** error log.

```
[student@serverX ~]$ sudo tail /var/log/httpd/error_log
```

- 8.2. That **"undefined function mysqli_connect"** error looks like a missing PHP library. Install the *php-mysql* package to install the needed bits.

```
[student@serverX ~]$ sudo yum install php-mysql
```

- 8.3. Restart **httpd.service** to make any running PHP instances aware of the new libraries.

```
[student@serverX ~]$ sudo systemctl restart httpd.service
```

- 9. From **desktopX**, use **firefox** to verify that the application now works. When a database listing is returned, execute **lab phpdb grade** on **desktopX** to verify your work.

- 9.1.

```
[student@desktopX ~]$ lab phpdb grade
```

Lab: Providing Apache HTTPD Web Service

Performance checklist

In this lab, you will configure your **serverX** to serve a Python WSGI web application over HTTPS.

Resources:	
Files:	/home/student/webapp.wsgi
Machines:	desktopX and serverX

Outcomes:

A TLS virtual host serving a Python WSGI web application on the **https://webappX.example.com** domain.

Before you begin...

- Reset your **desktopX** system.
- Reset your **serverX** system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab webapp setup
```

Your developers are working on a new web-based application written in Python. They have asked you to configure the **httpd** service on **serverX** to run this application on the **https://webappX.example.com** domain, using TLS encryption.

The application is programmed to be run behind a web server that support **WSGI**.

To configure the application, you will need the following information:

File	Download location
TLS certificate	http://classroom/pub/tls/certs/webappX.crt
TLS private key	http://classroom/pub/tls/private/webappX.key
TLS CA certificate	http://classroom/pub/example-ca.crt
Python application	/home/student/webapp.wsgi

The web application should be copied to a suitable location outside of any **DocumentRoot**.

Once you are done with your work, you can run the following command on your **desktopX** machine to validate your work:

```
[student@desktopX ~]$ lab webapp grade
```

1. Install the packages needed for **httpd**, **TLS**, and **WSGI** support.
2. Create a suitable location on the file system on **serverX** to host the web application, then copy the application there.

3. Download all the certificates and keys that you will need to configure TLS.
4. Configure a new TLS name-based virtual host for `https://webappX.example.com`, using the web application for `/`, and using the keys and certificates you downloaded earlier.
5. Start and enable your web server, and make the TLS-enabled web server available to the outside.
6. From your **desktopX** system, test the new web application.
7. Validate your work by running **lab webapp grade** on **desktopX**.

Solution

In this lab, you will configure your **serverX** to serve a Python WSGI web application over HTTPS.

Resources:	
Files:	<code>/home/student/webapp.wsgi</code>
Machines:	desktopX and serverX

Outcomes:

A TLS virtual host serving a Python WSGI web application on the **https://webappX.example.com** domain.

Before you begin...

- Reset your **desktopX** system.
- Reset your **serverX** system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab webapp setup
```

Your developers are working on a new web-based application written in Python. They have asked you to configure the **httpd** service on **serverX** to run this application on the **https://webappX.example.com** domain, using TLS encryption.

The application is programmed to be run behind a web server that support **WSGI**.

To configure the application, you will need the following information:

File	Download location
TLS certificate	<code>http://classroom/pub/tls/certs/webappX.crt</code>
TLS private key	<code>http://classroom/pub/tls/private/webappX.key</code>
TLS CA certificate	<code>http://classroom/pub/example-ca.crt</code>
Python application	<code>/home/student/webapp.wsgi</code>

The web application should be copied to a suitable location outside of any **DocumentRoot**.

Once you are done with your work, you can run the following command on your **desktopX** machine to validate your work:

```
[student@desktopX ~]$ lab webapp grade
```

1. Install the packages needed for **httpd**, **TLS**, and **WSGI** support.

- 1.1.

```
[student@serverX ~]$ sudo yum install httpd mod_ssl mod_wsgi
```

2. Create a suitable location on the file system on **serverX** to host the web application, then copy the application there.

- 2.1. Create a new directory `/srv/webappX/www`.


```
[student@serverX ~]$ sudo mkdir -p /srv/webappX/www
```

- 2.2. Copy the application to its new home.

```
[student@serverX ~]$ sudo cp ~/webapp.wsgi /srv/webappX/www/
```

- 2.3. Reset the SELinux context on the new directory.

```
[student@serverX ~]$ sudo restorecon -Rv /srv/webappX
```

3. Download all the certificates and keys that you will need to configure TLS.

- 3.1. Download the **example.com** CA certificate to **/etc/pki/tls/certs**.

```
[student@serverX ~]$ cd /etc/pki/tls/certs
[student@serverX certs]$ sudo wget http://classroom.example.com/pub/example-
ca.crt
```

- 3.2. Download the **webappX.example.com** certificate to **/etc/pki/tls/certs**.

```
[student@serverX certs]$ sudo wget http://classroom.example.com/pub/tls/certs/
webappX.crt
```

- 3.3. Download the **webappX** private key to **/etc/pki/tls/private/** and set the permissions to **0600**.

```
[student@serverX certs]$ cd /etc/pki/tls/private
[student@serverX private]$ sudo wget http://classroom.example.com/pub/tls/
private/webappX.key
[student@serverX private]$ sudo chmod 0600 webappX.key
```

4. Configure a new TLS name-based virtual host for **https://webappX.example.com**, using the web application for **/**, and using the keys and certificates you downloaded earlier.

- 4.1. Create a new file **/etc/httpd/conf.d/webappX.conf** with the following content:

```
<VirtualHost *:443>
  ServerName webappX.example.com
  SSLEngine on
  SSLProtocol all -SSLv2 -SSLv3
  SSLCipherSuite HIGH:MEDIUM:!aNULL:!MD5
  SSLHonorCipherOrder on
  SSLCertificateFile /etc/pki/tls/certs/webappX.crt
  SSLCertificateKeyFile /etc/pki/tls/private/X.key
  SSLCertificateChainFile /etc/pki/tls/certs/example-ca.crt
  WSGIScriptAlias / /srv/webappX/www/webapp.wsgi
</VirtualHost>
```

- 4.2. Add a **<Directory>** block to your new configuration to allow access to the web application:

```
<Directory /srv/webappX/www>  
  require all granted  
</Directory>
```

5. Start and enable your web server, and make the TLS-enabled web server available to the outside.

- 5.1. Start and enable **httpd**.

```
[student@serverX ~]$ sudo systemctl start httpd.service  
[student@serverX ~]$ sudo systemctl enable httpd.service
```

- 5.2. Open **https** in the default zone of the firewall on **serverX**.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --add-service=https  
[student@serverX ~]$ sudo firewall-cmd --reload
```

6. From your **desktopX** system, test the new web application.

- 6.1. Point an instance of **firefox** at `https://webappX.example.com`. Either import the **example.com** CA certificate as a trusted CA first, or accept the certificate warning.

- 6.2. Alternatively, you can run the following command on **desktopX**:

```
[student@desktopX ~]$ curl -k https://webappX.example.com
```

7. Validate your work by running **lab webapp grade** on **desktopX**.

- 7.1.

```
[student@desktopX ~]$ lab webapp grade
```

Summary

Configuring Apache HTTPD

In this section, students learned how to identify the key configuration files, log files, and content directories used by Apache **httpd**

Configuring and Troubleshooting Virtual Hosts

In this section, students learned how to configure Apache **httpd** to provide IP-based and name-based virtual hosts.

Configuring HTTPS

In this section, students learned how to configure Apache **httpd** to provide TLS-encrypted virtual hosts.

Integrating Dynamic Web Content

In this section, students learned how to configure Apache **httpd** to serve dynamic database-driven web content.



CHAPTER 11

WRITING BASH SCRIPTS

Overview	
Goal	To write simple, well-structured shell scripts using Bash's shell expansion features and for-loop construct.
Objectives	<ul style="list-style-type: none">• To write simple shell scripts using Bash.
Sections	<ul style="list-style-type: none">• Bash Shell Scripting Basics (and Practice)
Lab	<ul style="list-style-type: none">• Writing Bash Scripts

Bash Shell Scripting Basics

Objectives

After completing this section, students should be able to write and debug simple Bash shell scripts which utilize variable expansion, command line substitution, arithmetic expansion, and for loops.

Bash scripting basics

Many simple day-to-day system administration tasks can be accomplished by the numerous Linux command-line tools available to administrators. However, tasks with greater complexity often require the chaining together of multiple commands. In these situations, Linux command-line tools can be combined with the offerings of the Bash shell to create powerful shell scripts to solve real-world problems.

In its simplest form, a Bash shell script is simply an executable file composed of a list of commands. However, when well-written, a shell script can itself become a powerful command-line tool when executed on its own, and can even be further leveraged by other scripts.

Proficiency in shell scripting is essential to the success of Linux system administrators in all operational environments. Working knowledge of shell scripting is especially crucial in enterprise environments, where its use can translate to improved efficiency and accuracy of routine task completion.

Choosing a programming language

While Bash shell scripting can be used to accomplish many tasks, it may not be the right tool for all scenarios. Administrators have a wide variety of programming languages at their disposal, such as C, C++, Perl, Python, Ruby and other programming languages. Each programming language has its strengths and weaknesses, and as such, none of the programming languages are the right tool for every situation.

Bash shell scripts are a good choice for tasks which can be accomplished mainly by calling other command-line utilities. If the task involves heavy data processing and manipulation, other languages such as Perl or Python will be better suited for the job. While Bash supports arithmetic operations, they are limited to simple integer arithmetic. For more complex arithmetic operations, C or C++ should be considered. If a solution requires the use of arrays, Bash is probably not the best tool. Bash has supported one-dimensional arrays for some time, and the latest version even supports associative arrays. However, Perl or Python have much better array functionality, with the ability to accommodate multidimensional arrays.

As administrators become proficient at shell scripting, they will gain more knowledge regarding its capabilities and limitations. This experience combined with exposure to other programming languages over time, will provide administrators with a better understanding of the advantages and disadvantages of each, and which problems each one is best suited for.

Creating and executing Bash shell scripts

A Bash shell script can be created by opening a new empty file in a text editor. While any text editor can be used, advanced editors, such as *vim* or *emacs*, understand Bash shell syntax and can provide color-coded highlighting. This highlighting can be a tremendous help for spotting syntactical errors, such as unpaired quotes, unclosed brackets, and other common blunders.

The command interpreter

The first line of a Bash shell script begins with '#!', also commonly referred to as a *sharp-bang*, or the abbreviated version, *sha-bang*. This two-byte notation is technically referred to as the *magic pattern*. It indicates that the file is an executable shell script. The path name that follows is the *command interpreter*, the program that should be used to execute the script. Since Bash shell scripts are to be interpreted by the Bash shell, they begin with the following first line.

```
#!/bin/bash
...
```

Executing a Bash shell script

After a Bash shell script is written, its file permissions and ownership need to be modified so that it is executable. Execute permission is modified with the **chmod** command, possibly in conjunction with the **chown** command to change the file ownership of the script accordingly. Execute permission should only be granted to the users that the script is intended for.

Once a Bash shell script is executable, it can be invoked by entering its name on the command line. If only the base name of the script file is entered, Bash will search through the directories specified in the shell's **PATH** environmental variable, looking for the first instance of an executable file matching that name. Administrators should avoid script names that match other executable files, and should also ensure that the **PATH** variable is correctly configured on their system so that the script will be the first match found by the shell. The **which** command, followed by the file-name of the executable script displays in which directory the script resides that is executed when the script name is invoked as a command resides.

```
[student@serverX ~]$ which hello
~/bin/hello

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

Displaying output

The **echo** command can be used to display arbitrary text by passing the text as an argument to the command. By default, the text is directed to *standard out (STDOUT)*, but can also be directed to *standard error (STDERR)* using output redirection. In the following simple Bash script, the **echo** command displays the message "Hello, world" to *STDOUT*.

```
[student@serverX ~]$ cat hello
#!/bin/bash

echo "Hello, world"

[student@serverX ~]$ ./hello
Hello, world
```

While seemingly simple in its function, the **echo** command is widely used in shell scripts due to its usefulness for various purposes. It is commonly used to display informational or error messages during the script execution. These messages can be a helpful indicator of the progress of a script and can be directed either to standard out, standard error, or be redirected to a log file for archiving. When displaying error messages, it is good practice to direct them to *STDERR* to make it easier to differentiate error messages from normal status messages.

```
[student@serverX ~]$ cat hello
#!/bin/bash

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

[student@serverX ~]$ ./hello 2> hello.log
Hello, world

[student@serverX ~]$ cat hello.log
ERROR: Houston, we have a problem.
```

The **echo** command can also be very helpful when trying to debug a problematic shell script. The addition of **echo** statements to the portion of the script that is not behaving as expected can help clarify the commands being executed, as well as the values of variables being invoked.

Quoting special characters

A number of characters or words have special meanings to the Bash shell in specific context. There are situations when the literal values, rather than the special meanings, of these characters or words are desired. For example, the **#** character is interpreted by Bash as the beginning of a comment and is therefore ignored, along with everything following it on the same line. If this special meaning is not desired, then Bash needs to be informed that the **#** character is to be treated as a literal value. The meanings of special characters or words can be disabled through the use of the escape character, ****, single quotes, **' '**, or double quotes, **" "**.

The escape character, ****, removes the special meaning of the single character immediately following it. For example, to display the literal string **# test** with the **echo** command, the **#** character must not be interpreted by Bash with special meaning. The escape character can be placed in front of the **#** character to disable its special meaning.

```
[student@serverX ~]$ echo # not a comment

[student@serverX ~]$ echo \# not a comment
# not a comment
```

The escape character, ****, only removes the special meaning of a single character. When more than one character in a text string needs to be escaped, users can either use the escape character multiple times or employ single quotes, **' '**. Single quotes preserve the literal meaning of all characters they enclose. The following example demonstrates how single quotes can be used when multiple characters need to be escaped.

```
[student@serverX ~]$ echo # not a comment #

[student@serverX ~]$ echo \# not a comment #
# not a comment

[student@serverX ~]$ echo \# not a comment \#
# not a comment #

[student@serverX ~]$ echo '# not a comment #'
# not a comment #
```

While single quotes preserve the literal value of all characters they enclose, double quotes differ in that they do not preserve the literal value of the dollar sign, **\$**, the back-ticks, **`**, and the backslash, ****. When enclosed with double quotes, the dollar sign and back-ticks preserve their

special meaning, and the special meaning of the backslash character is only retained when it precedes a dollar sign, back-tick, double quote, backslash, or newline.

```
[student@serverX ~]$ echo '$HOME'
$HOME

[student@serverX ~]$ echo '`pwd`'
`pwd`

[student@serverX ~]$ echo '"Hello, world"'
"Hello, world"

[student@serverX ~]$ echo "$HOME"
/home/student

[student@serverX ~]$ echo "`pwd`"
/home/student

[student@serverX ~]$ echo "'Hello, world'"
Hello, world

[student@serverX ~]$ echo "\$HOME"
$HOME

[student@serverX ~]$ echo "`\`pwd\`"
`pwd`

[student@serverX ~]$ echo "\"Hello, world\""
"Hello, world"
```

Using variables

As the complexity of a shell script increases, it is often helpful to make use of variables. A variable serves as a container, within which a shell script can store data in memory. Variables make it easy to access and modify the stored data during a script's execution.

Assigning values to variables

Data is assigned as a value to a variable via the following syntax:

```
VARIABLENAME=value
```

While variable names are typically uppercase letters, they can be made up of numbers, letters (uppercase and lowercase), and the underscore character, '_'. However, a variable name cannot start with a number. The equal sign, '=', is used to assign values to variables and must not be separated from the variable name or the value by spaces. The following are some examples of valid variable names.

```
COUNT=40
first_name=John
file1=/tmp/abc
_ID=RH123
```

Two common types of data stored in variables are integer values and string values. It is good practice to quote string values when assigning them to variables, since the space character is interpreted by Bash as a word separator when not enclosed within single or double quotes. Whether single or double quotes should be used to enclose variable values depends on how characters with special meanings to Bash should be treated.


```
full_name='John Doe'
full_name="$FIRST $LAST"
price='$1'
```

Expanding variable values

The value of a variable can be recalled through a process known as *variable expansion* by preceding the variable name with a dollar sign, **\$**. For example, the value of the **VARIABLENAME** variable can be referenced with **\$VARIABLENAME**. The **\$VARIABLENAME** syntax is the simplified version of the brace-quoted form of variable expansion, **\${VARIABLENAME}**. While the simplified form is usually acceptable, there are situations where the brace-quoted form must be used to remove ambiguity and avoid unexpected results.

In the following example, without the use of brace quotes, Bash will interpret **\$FIRST_\$LAST** as the variable **\$FIRST_** followed by the variable **\$LAST**, rather than the variables **\$FIRST** and **\$LAST** separated by the '_' character. Therefore, brace quoting must be used for variable expansion to function properly in this scenario.

```
[student@serverX ~]$ FIRST=Jane
[student@serverX ~]$ FIRST=John
[student@serverX ~]$ LAST=Doe
[student@serverX ~]$ echo $FIRST_$LAST
JaneDoe
[student@serverX ~]$ echo ${FIRST}_$LAST
John_Doe
```

Using Bash shell expansion features

Aside from variable expansion, the Bash shell offers several other types of shell expansion features. Of these, command substitution and arithmetic expansion can be useful in Bash shell scripting, and are commonly used.

Command substitution

Command substitution replaces the invocation of a command with the output from its execution. This feature allows the output of a command to be used in a new context, such as the argument to another command, the value for a variable, and the list for a loop construct.

Command substitution can be invoked with the old form of enclosing the command in back-ticks, such as **`<COMMAND>`**. However, the preferred method is to use the newer **\$()** syntax, **\$(<COMMAND>)**.

```
[student@serverX ~]$ echo "Current time: `date`"
Current time is Thu Jun  5 16:24:24 EDT 2014.
[student@serverX ~]$ echo "Current time: $(date)"
Current time is Thu Jun  5 16:24:30 EDT 2014.
```

The newer syntax is preferred since it allows for nesting of command substitutions. In the following nested command substitution example, the output of the **find** command is used as arguments for the **tar** command, which then has its output stored into the variable **TAROUTPUT**.

```
[root@serverX ~]# TAROUTPUT=$(tar cvf /tmp/incremental_backup.tar $(find /etc -type f -
mtime -1))

[root@serverX ~]# echo $TAROUTPUT
/etc/group /etc/gshadow /etc/shadow- /etc/passwd /etc/shadow /etc/passwd- /etc/tuned/
active_profile /etc/rht /etc/group- /etc/gshadow- /etc/resolv.conf
```

Arithmetic expansion

Bash's arithmetic expansion can be used to perform simple integer arithmetic operations, and uses the syntax `$(<EXPRESSION>)`. When enclosed within `$([])`, arithmetic expressions are evaluated by Bash and then replaced with their results. Bash performs variable expansion and command substitution on the enclosed expression before its evaluation. Like command line substitution, nesting of arithmetic substitutions is allowed.

```
[student@serverX ~]$ echo ${1+1}
2

[student@serverX ~]$ echo ${2*2}
4

[student@serverX ~]$ COUNT=1; echo ${${COUNT+1}*2}
4
```

Space characters are allowed in the expression used within an arithmetic expansion. The use of space characters can improve readability in complicated expressions or when variables are included.

```
[student@serverX ~]$ SEC_PER_MIN=60

[student@serverX ~]$ MIN_PER_HR=60

[student@serverX ~]$ HR_PER_DAY=24

[student@serverX ~]$ SEC_PER_DAY=${ SEC_PER_MIN * $MIN_PER_HR * $HR_PER_DAY }

[student@serverX ~]$ echo "There are $SEC_PER_DAY seconds in a day."
There are 86400 seconds in a day.
```

The following are some of the commonly used operators in arithmetic expressions, along with their meanings.

Operator	Meaning
<VARIABLE>++	variable post-increment
<VARIABLE>--	variable post-decrement
++<VARIABLE>	variable pre-increment
--<VARIABLE>	variable pre-decrement
-	unary minus
+	unary plus
**	exponentiation
*	multiplication
/	division

Operator	Meaning
%	remainder
+	addition
-	subtraction

When multiple operators exist in an expression, Bash will evaluate certain operators in order according to their precedence. For example, multiplication and division operators have a higher precedence than addition and subtraction. Parentheses can be used to group sub expressions if the evaluation order desired differs from the default precedence.

```
[student@serverX ~]$ echo $[ 1 + 1 * 2 ]
3
[student@serverX ~]$ echo $[ (1 + 1) * 2 ]
4
```

The following table lists the order of precedence for commonly used arithmetic operators from highest to lowest. Operators that have equal precedence are listed together.

Operator	Meaning
<VARIABLE>++, <VARIABLE>--	variable post-increment and post-decrement
++<VARIABLE>, --<VARIABLE>	variable pre-increment and pre-decrement
-, +	unary minus and plus
**	exponentiation
*, /, %	multiplication, division, remainder
+, -	addition, subtraction

Iterating with the for loop

System administrators often encounter repetitive tasks in their day-to-day activities. Repetitive tasks can take the form of executing an action multiple times on an target, such as checking a process every minute for 10 minutes to see if it has completed. Task repetition can also take the form of executing an action a single time across multiple targets, such as performing a database backup of each database on a system. The *for loop* is one of the multiple shell looping constructs offered by Bash, and can be used for task iterations.

Using the for loop

Bash's for-loop construct uses the following syntax. The loop processes the items provided in **<LIST>** in order one by one and exits after processing the last item on the list. Each item in the list is temporarily stored as the value of **<VARIABLE>**, while the for loop executes the block of commands contained in its construct. The naming of the variable is arbitrary. Typically, the variable value is referenced by commands in the command block.

```
for <VARIABLE> in <LIST>; do
    <COMMAND>
    ...
    <COMMAND> referencing <VARIABLE>
done
```

The list of items provided to a for loop can be supplied in several ways. It can be a list of items entered directly by the user, or be generated from different types of shell expansion, such as variable expansion, brace expansion, file name expansion, and command substitution. Some examples that demonstrate the different ways lists can be provided to for loops follow.

```
[student@serverX ~]$ for HOST in host1 host2 host3; do echo $HOST; done
host1
host2
host3

[student@serverX ~]$ for HOST in host{1,2,3}; do echo $HOST; done
host1
host2
host3

[student@serverX ~]$ for HOST in host{1..3}; do echo $HOST; done
host1
host2
host3

[student@serverX ~]$ for FILE in file*; do ls $FILE; done
filea
fileb
filec

[student@serverX ~]$ for FILE in file{a..c}; do ls $FILE; done
filea
fileb
filec

[student@serverX ~]$ for PACKAGE in $(rpm -qa | grep kernel); do echo "$PACKAGE was
  installed on $(date -d @$(rpm -q --qf "%{INSTALLTIME}\n" $PACKAGE))"; done
abrt-addon-kerneloops-2.1.11-12.el7.x86_64 was installed on Tue Apr 22 00:09:07 EDT 2014
kernel-3.10.0-121.el7.x86_64 was installed on Thu Apr 10 15:27:52 EDT 2014
kernel-tools-3.10.0-121.el7.x86_64 was installed on Thu Apr 10 15:28:01 EDT 2014
kernel-tools-libs-3.10.0-121.el7.x86_64 was installed on Thu Apr 10 15:26:22 EDT 2014

[student@serverX ~]$ for EVEN in $(seq 2 2 8); do echo "$EVEN"; done; echo "Who do we
  appreciate?"
2
4
6
8
Who do we appreciate?
```

Troubleshooting shell script bugs

Inevitably, administrators who write, use, or maintain shell scripts will encounter bugs with a script. Bugs are typically due to typographical errors, syntactical errors, or poor script logic.

A good way to deal with shell scripting bugs is to make a concerted effort to prevent them from occurring in the first place during the authoring of the script. As previously mentioned, using a text editor with Bash syntactical highlighting can help make mistakes more obvious when writing scripts. Another easy way to avoid introducing bugs into scripts is by adhering to good practices during the creation of the script.

Good styling practices

Use comments to help clarify to readers the purpose and logic of the script. The top of every script should include comments providing an overview of the script's purpose, intended actions,

and general logic. Also use comments throughout the script to clarify the key portions, and especially sections that may cause confusion. Comments will not only aid other users in the reading and debugging of the script, but will also often help the author recall the workings of the script once some time has passed.

Structure the contents of the script to improve readability. As long as the syntax is correct, the command interpreter will flawlessly execute the commands within a script with absolutely no regard for their structure or formatting. Here are some good practices to follow:

- Break up long commands into multiple lines of smaller code chunks. Shorter pieces of code are much easier for readers to digest and comprehend.
- Line up the beginning and ending of multiline statements to make it easier to see that control structures begin and end, and whether they are being closed properly.
- Indent lines with multiline statements to represent the hierarchy of code logic and the flow of control structures.
- Use line spacing to separate command blocks to clarify when one code section ends and another begins.
- Use consistent formatting through the entirety of a script.

When utilized, these simple practices can make it significantly easier to spot mistakes during authoring, as well as improve the readability of the script for future readers. The following example demonstrates how the incorporation of comments and spacing can greatly improve script readability.

```
#!/bin/bash
for PACKAGE in $(rpm -qa | grep kernel); do echo "$PACKAGE was installed on $(date -d @
$(rpm -q --qf "%{INSTALLTIME}\n" $PACKAGE))"; done
```

```
#!/bin/bash
#
# This script provides information regarding when kernel-related packages
# are installed on a system by querying information from the RPM database.
#
# Variables
PACKAGETYPE=kernel
PACKAGES=$(rpm -qa | grep $PACKAGETYPE)
# Loop through packages
for PACKAGE in $PACKAGES; do
    # Determine package install date and time
    INSTALLEPOCH=$(rpm -q --qf "%{INSTALLTIME}\n" $PACKAGE)
    # RPM reports time in epoch, so need to convert
    # it to date and time format with date command
    INSTALLDATETIME=$(date -d @$INSTALLEPOCH)
    # Print message
    echo "$PACKAGE was installed on $INSTALLDATETIME"
done
```

Do not make assumptions regarding the outcome of actions taken by a script. This is especially true of inputs to the script, such as command-line arguments, input from users, command

substitutions, variable expansions, and file name expansions. Rather than making assumptions about the integrity of these inputs, make the worthwhile effort to employ the use of proper quoting and sanity checking.

The same caution should be utilized when acting upon entities external to the script. This includes interacting with files, and calling external commands. Make use of Bash's vast number of file and directory tests when interacting with files and directories. Perform error checking on the exit status of commands rather than counting on their success and blindly continuing along with the script when an unexpected error occurs.

The extra steps taken to rule out assumptions will increase the script's robustness, and keep it from being easily derailed and then inflicting unintended and unnecessary damage to a system. A couple of seemingly harmless lines of code, such as the ones that follow, make very risky assumptions about command execution outcome and file name expansion. If the directory change fails, either due to directory permissions or the directory being nonexistent, the subsequent file removal will be performed on a list of unknown files in an unintended directory.

```
cd $TMPDIR
rm *
```

Lastly, while well-intentioned administrators may employ good practices when authoring their scripts, not all will always agree on what constitutes good practices. Administrators should do themselves and others a favor and always apply their practices consistently through the entirety of their scripts. They should also be considerate and understanding of individual differences when it comes to programming styles and formatting in scripts authored by others. When modifying others' scripts, administrators should follow the existing structure, formatting, and practices used by the original author, rather than imposing their own style on a portion of the script and destroying the script's consistency, and thereby ruining its readability and future maintainability.

Debug and verbose modes

If despite best efforts, bugs are introduced into a script, administrators will find Bash's debug mode extremely useful. To activate the debug mode on a script, add the `-x` option to the command interpreter in the first line of the script.

```
#!/bin/bash -x
```

Another way to run a script in debug mode is to execute the script as an argument to Bash with the `-x` option.

```
[student@serverX bin]$ bash -x <SCRIPTNAME>
```

Bash's debug mode will print out commands executed by the script prior to their execution. The results of all shell expansion performed will be displayed in the printout. The following example shows the extra output that is displayed when debug mode is activated.

```
[student@serverX bin]$ cat filesize
#!/bin/bash

DIR=/home/student/tmp

for FILE in $DIR/*; do
  echo "File $FILE is $(stat --printf='%s' $FILE) bytes."
done

[student@serverX bin]$ ./filesize
File /home/student/tmp/filea is 133 bytes.
File /home/student/tmp/fileb is 266 bytes.
File /home/student/tmp/filec is 399 bytes.

[student@serverX bin]$ bash -x ./filesize
+ DIR=/home/student/tmp
+ for FILE in '$DIR/*'
++ stat --printf=%s /home/student/tmp/filea
+ echo 'File /home/student/tmp/filea is 133 bytes.'
File /home/student/tmp/filea is 133 bytes.
+ for FILE in '$DIR/*'
++ stat --printf=%s /home/student/tmp/fileb
+ echo 'File /home/student/tmp/fileb is 266 bytes.'
File /home/student/tmp/fileb is 266 bytes.
+ for FILE in '$DIR/*'
++ stat --printf=%s /home/student/tmp/filec
+ echo 'File /home/student/tmp/filec is 399 bytes.'
File /home/student/tmp/filec is 399 bytes.
```

While Bash's debug mode provides helpful information, the voluminous output may actually become more hindrance than help for troubleshooting, especially as the lengths of scripts increase. Fortunately, the debug mode can be enabled partially on just a portion of a script, rather than on its entirety. This feature is especially useful when debugging a long script and the source of the problem has been narrowed to a portion of the script.

Debugging can be turned on at a specific point in a script by inserting the command **set -x** and turned off by inserting the command **set +x**. The following demonstration shows the previous example script with debugging enabled just for the command line enclosed in the for loop.

```
[student@serverX bin]$ cat filesize
#!/bin/bash

DIR=/home/student/tmp

for FILE in $DIR/*; do
    set -x
    echo "File $FILE is $(stat --printf='%s' $FILE) bytes."
    set +x
done

[student@serverX bin]$ ./filesize
++ stat --printf=%s /home/student/tmp/filea
+ echo 'File /home/student/tmp/filea is 133 bytes.'
File /home/student/tmp/filea is 133 bytes.
+ set +x
++ stat --printf=%s /home/student/tmp/fileb
+ echo 'File /home/student/tmp/fileb is 266 bytes.'
File /home/student/tmp/fileb is 266 bytes.
+ set +x
++ stat --printf=%s /home/student/tmp/filec
+ echo 'File /home/student/tmp/filec is 399 bytes.'
File /home/student/tmp/filec is 399 bytes.
+ set +x
```

In addition to debug mode, Bash also offers a verbose mode, which can be invoked with the **-v** option. In verbose mode, Bash will print each command to standard out prior to its execution.

```
[student@serverX bin]$ cat filesize
#!/bin/bash

DIR=/home/student/tmp

for FILE in $DIR/*; do
    echo "File $FILE is $(stat --printf='%s' $FILE) bytes."
done

[student@serverX bin]$ bash -v ./filesize
stat --printf='%s' $FILE) bytes."
stat --printf='%s' $FILE) bytes.
stat --printf='%s' $FILE
File /home/student/tmp/filea is 133 bytes.
stat --printf='%s' $FILE) bytes."
stat --printf='%s' $FILE) bytes.
stat --printf='%s' $FILE
File /home/student/tmp/fileb is 266 bytes.
stat --printf='%s' $FILE) bytes."
stat --printf='%s' $FILE) bytes.
stat --printf='%s' $FILE
File /home/student/tmp/filec is 399 bytes.
```

Like the debug feature, the verbose feature can also be turned on and off at specific points in a script by inserting the **set -v** and **set +v** lines, respectively.



References

bash(1), **magic**(5), **echo**(1), **echo**(1p), and **seq**(1) man pages

Practice: Writing Bash Scripts

Guided exercise

In this lab, you will create a Bash shell script to automate the process of individually backing up every MariaDB database on **serverX** and generating a report providing statistics on each database backup.

Resources:	
Machines:	serverX

Outcomes:

A Bash script which automates the process of performing individual backups of each database in MariaDB. The script will also report the size statistics on each database backup.

Before you begin...

- Reset **serverX**.
- Log into **serverX** and become root with **sudo -i**.

You have been tasked with writing a Bash shell script to perform a backup of all MariaDB databases on a server using **mysqldump**. Each database will be backed up to the directory **/dbbackup** and will be named **DATABASENAME.dump**. Your script will print out the message **'Backing up "DATABASENAME"'** as it initiates the dump of each database.

Once the backups for all the databases are completed, your script will generate a report showing the name of each database backup, its size, and the percentage of the total database dump size it accounts for. The data for each database should be lined up in columns for readability.

The complete output of your script should look like the following:

```
[root@serverX ~]# /usr/local/sbin/dbbackup
Backing up "mysql"
Backing up "test"

/dbbackup/mysql.dump          514664 99%
/dbbackup/test.dump           1261 0%
```

- 1. Install the *mariadb-server* package, then enable and start the **mariadb** service.
 - 1.1. Install *mariadb-server* package with **yum**.

```
[root@serverX ~]# yum install -y mariadb-server
```

- 1.2. Enable and start **mariadb**.

```
[root@serverX ~]# systemctl enable mariadb
ln -s '/usr/lib/systemd/system/mariadb.service' '/etc/systemd/system/
multi-user.target.wants/mariadb.service'
[root@serverX ~]# systemctl start mariadb
```

- 2. Create the database backup directory.

```
[root@serverX ~]# mkdir /dbbackup
```

- 3. On the command line, formulate a command to generate a list of database names, excluding the system databases **information_schema** and **performance_schema**.

- 3.1. Issue the **SHOW DATABASES** command to **mysql** as the **root** MySQL user.

```
[root@serverX ~]# mysql -u root -e 'SHOW DATABASES'
+-----+
| Database           |
+-----+
| information_schema |
| mysql              |
| performance_schema |
| test               |
+-----+
```

- 3.2. Utilize the **--skip-column-names** and **-E** formatting options to simplify the output for parsing.

```
[root@serverX ~]# mysql --skip-column-names -E -u root -e 'SHOW DATABASES'
***** 1. row *****
information_schema
***** 2. row *****
mysql
***** 3. row *****
performance_schema
***** 4. row *****
test
```

- 3.3. Exclude the row header lines and the two system databases from the output.

```
[root@serverX ~]# mysql --skip-column-names -E -u root -e 'SHOW
DATABASES' | grep -v '^*' | grep -v '^information_schema$' | grep -v
'^performance_schema$'
mysql
test
```

- 4. Create your script. Store the MySQL user, the formatting options, the **SHOW DATABASES** command, and the backup directory as variables.

- 4.1. Create the new script file with a text editor.

```
[root@serverX ~]# vim /usr/local/sbin/dbbackup
```

- 4.2. Specify the interpreter program for the script.

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

- 4.3. Set the variables.

```
# Variables
DBUSER=root
FMTOPTIONS='--skip-column-names -E'
COMMAND='SHOW DATABASES'
BACKUPDIR=/dbbackup
```

- 5. Initiate a for loop and loop through the list of databases to back up each one to the **/dbbackup** directory.

- 5.1. Initiate the for loop by passing in a list of database names via command substitution.

```
# Backup non-system databases
for DBNAME in $(mysql $FMTOPTIONS -u $DBUSER -e "$COMMAND" | grep -v ^* |
grep -v information_schema | grep -v performance_schema); do
```

- 5.2. Add the commands to be executed within each loop.

```
echo "Backing up \"$DBNAME\""
mysqldump -u $DBUSER $DBNAME > $BACKUPDIR/$DBNAME.dump
```

- 5.3. Close the for loop.

```
done
```

- 6. Generate a report of each database's name, dump size, and the percentage of the total dump size it accounts for.

- 6.1. Initiate a for loop to iterate through and total up the size of each database dump in the **/dbbackup** directory.

```
# Add up size of all database dumps
for DBDUMP in $BACKUPDIR/*; do
```

- 6.2. Add the commands to be executed within each loop.

```
SIZE=$(stat --printf "%s\n" $DBDUMP)
TOTAL=$(( $TOTAL + $SIZE )
```

- 6.3. Close the for loop.

```
done
```

- 6.4. Create a for loop to iterate through and report on each database dump.

```
# Report name, size, and percentage of total for each database dump
echo
for DBDUMP in $BACKUPDIR/*; do
```

- 6.5. Add the commands to be executed within each loop.

```
SIZE=$(stat --printf "%s\n" $DBDUMP)
echo "$DBDUMP,$SIZE,$[ 100 * $SIZE / $TOTAL ]%"
```

- 6.6. Close the for loop.

```
done
```

- 7. Save and execute the script.

- 7.1. Make the script executable.

```
[root@serverX ~]# chmod u+x /usr/local/sbin/dbbackup
```

- 7.2. Execute the script.

```
[root@serverX ~]# /usr/local/sbin/dbbackup
Backing up "mysql"
Backing up "test"

/dbbackup/mysql.dump,514664,99%
/dbbackup/test.dump,1261,0%
```

Lab: Writing Bash Scripts

Performance checklist

In this lab, you will write a simple Bash script to automate the creation of user accounts.

Resources:	
Files:	/tmp/support/newusers
Machines:	serverX

Outcomes:

A shell script, `/usr/local/sbin/mkaccounts`, which creates new user accounts based on information provided in a data file.

Before you begin...

- Reset **serverX**.
- Log into **serverX** and become the **root** user with **sudo -i**.
- Run **lab bashbasic setup**.

Your company provides hosting service to customers, and you have been tasked with writing a Bash shell script called `/usr/local/sbin/mkaccounts` to automate the process of creating accounts for new customers. At the end of each day, a colon-separated data file called `/tmp/support/newusers` is created and contains information on new customers that have signed up. The script will read through this data file and create a user account for each new customer.

The sales department has requested that the script also generate a report breaking down the new customers by support tier, so sales staff can stay informed regarding trends of support tier purchases. For each support tier, they would like the report to detail the total number of new customers and the percentage of the day's new customers that chose the tier type.

Once you have completed the script, run **lab bashbasic grade** to verify its functionality.

1. Create your script. Set the file name, `/tmp/support/newusers`, as the value of the variable **NEWUSERSFILE**.
2. Loop through the entries in the **\$NEWUSERSFILE** data file and extract the values of the first, second, and last fields to obtain the first name, last name, and support tier level of each new customer. Create an account for each new customer such that the account name is the lowercase combination of their first initial and last name (i.e., jdoe) and the comment of the account is their full name (i.e., John Doe).
3. Create a report summarizing the number of new customers for each support tier and what percentage of the total new user accounts the support type comprises. The report should print as follows:

```
"Tier 1", "22", "44%"
"Tier 2", "15", "30%"
"Tier 3", "13", "26%"
```

4. Save and execute the script. Verify that the script performed as intended. If the script did not perform as intended, you will need to run **lab bashbasic clean** to purge out accounts that might have been created before rerunning your script.
5. Verify your script by running **lab bashbasic grade**.

Solution

In this lab, you will write a simple Bash script to automate the creation of user accounts.

Resources:	
Files:	/tmp/support/newusers
Machines:	serverX

Outcomes:

A shell script, `/usr/local/sbin/mkaccounts`, which creates new user accounts based on information provided in a data file.

Before you begin...

- Reset **serverX**.
- Log into **serverX** and become the **root** user with **sudo -i**.
- Run **lab bashbasic setup**.

Your company provides hosting service to customers, and you have been tasked with writing a Bash shell script called `/usr/local/sbin/mkaccounts` to automate the process of creating accounts for new customers. At the end of each day, a colon-separated data file called `/tmp/support/newusers` is created and contains information on new customers that have signed up. The script will read through this data file and create a user account for each new customer.

The sales department has requested that the script also generate a report breaking down the new customers by support tier, so sales staff can stay informed regarding trends of support tier purchases. For each support tier, they would like the report to detail the total number of new customers and the percentage of the day's new customers that chose the tier type.

Once you have completed the script, run **lab bashbasic grade** to verify its functionality.

1. Create your script. Set the file name, `/tmp/support/newusers`, as the value of the variable **NEWUSERSFILE**.
 - 1.1. Create the new script with a text editor.

```
[root@serverX ~]# vim /usr/local/sbin/mkaccounts
```

- 1.2. Specify the interpreter program for the script.

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

- 1.3. Set the **NEWUSERSFILE** variable.

```
# Variables
NEWUSERSFILE=/tmp/support/newusers
```

2. Loop through the entries in the **\$NEWUSERSFILE** data file and extract the values of the first, second, and last fields to obtain the first name, last name, and support tier level of each new customer. Create an account for each new customer such that the account name is the

lowercase combination of their first initial and last name (i.e., jdoe) and the comment of the account is their full name (i.e., John Doe).

- 2.1. Read through the entries in the file by initiating a for loop in combination with command substitution.

```
# Loop
for ENTRY in $(cat $NEWUSERSFILE); do
```

- 2.2. Extract and save the values of the first, second, and last fields.

```
# Extract first, last, and tier fields
FIRSTNAME=$(echo $ENTRY | cut -d: -f1)
LASTNAME=$(echo $ENTRY | cut -d: -f2)
TIER=$(echo $ENTRY | cut -d: -f4)
```

- 2.3. Compose the account name for the new customer.

```
# Make account name
FIRSTINITIAL=$(echo $FIRSTNAME | cut -c 1 | tr 'A-Z' 'a-z')
LOWERLASTNAME=$(echo $LASTNAME | tr 'A-Z' 'a-z')

ACCTNAME=$FIRSTINITIAL$LOWERLASTNAME
```

- 2.4. Create the account with the **useradd** command.

```
# Create account
useradd $ACCTNAME -c "$FIRSTNAME $LASTNAME"
```

- 2.5. Close the for loop.

```
done
```

3. Create a report summarizing the number of new customers for each support tier and what percentage of the total new user accounts the support type comprises. The report should print as follows:

```
"Tier 1", "22", "44%"
"Tier 2", "15", "30%"
"Tier 3", "13", "26%"
```

- 3.1. Determine the total number of new customers and store the value in a variable.

```
TOTAL=$(cat $NEWUSERSFILE | wc -l)
```

- 3.2. Determine the total number of new customers for each support tier and store the totals into variables.

```
TIER1COUNT=$(grep -c :1$ $NEWUSERSFILE)
TIER2COUNT=$(grep -c :2$ $NEWUSERSFILE)
TIER3COUNT=$(grep -c :3$ $NEWUSERSFILE)
```

- 3.3. Calculate the percentages for each support tier.

```
TIER1PCT=$(( TIER1COUNT * 100 / $TOTAL ))
TIER2PCT=$(( TIER2COUNT * 100 / $TOTAL ))
TIER3PCT=$(( TIER3COUNT * 100 / $TOTAL ))
```

- 3.4. Print the report.

```
# Print report
echo "\"Tier 1\", \"$TIER1COUNT\", \"$TIER1PCT%\""
echo "\"Tier 2\", \"$TIER2COUNT\", \"$TIER2PCT%\""
echo "\"Tier 3\", \"$TIER3COUNT\", \"$TIER3PCT%\""
```

4. Save and execute the script. Verify that the script performed as intended. If the script did not perform as intended, you will need to run **lab bashbasic clean** to purge out accounts that might have been created before rerunning your script.

- 4.1. Make the script executable.

```
[root@serverX ~]# chmod u+x /usr/local/sbin/mkaccounts
```

- 4.2. Execute the script.

```
[root@serverX ~]# /usr/local/sbin/mkaccounts
"Tier 1", "22", "44.00%"
"Tier 2", "15", "30.00%"
"Tier 3", "13", "26.00%"
```

5. Verify your script by running **lab bashbasic grade**.

Summary

Bash Shell Scripting Basics

In this section, students learned how to:

- Write, execute, and debug simple Bash shell scripts.
- Store and retrieve data during script execution with the use of variables.
- Enhance script functionality with command substitution and arithmetic expansion.
- Perform iterative execution of tasks with for loops.
- Troubleshoot Bash shell script issues with debug and verbose modes.



CHAPTER 12

BASH CONDITIONALS AND CONTROL STRUCTURES

Overview	
Goal	To use Bash conditionals and other control structures to write more sophisticated shell commands and scripts.
Objectives	<ul style="list-style-type: none">• Incorporate the use of positional parameters, exit status, test conditions, and conditional structures to implement flow control in Bash shell scripts.
Sections	<ul style="list-style-type: none">• Enhancing Bash Shell Scripts with Conditionals and Control Structures (and Practice)
Lab	<ul style="list-style-type: none">• Bash Conditionals and Control Structures

Enhancing Bash Shell Scripts with Conditionals and Control Structures

Objectives

After completing this section, students should be able to incorporate the use of positional parameters, exit status, test conditions, and conditional structures to implement flow control in Bash shell scripts.

Using Bash special variables

While user-defined variables provide a means for script authors to create containers to store values used by a script, Bash also provides some predefined variables, which can be useful when writing shell scripts. One type of predefined variable is positional parameters.

Positional parameters

Positional parameters are variables which store the values of command-line arguments to a script. The variables are named numerically. The variable **0** refers to the script name itself. Following that, the variable **1** is predefined with the first argument to the script as its value, the variable **2** contains the second argument, and so on. The values can be referenced with the syntax **\$1**, **\$2**, etc.



Important

While rare, when referencing values past the ninth positional parameter, the brace-quoted form of variable expansion must be used. For example, the value of the tenth positional argument must be referenced with the syntax **\${10}** rather than **\$10**. Otherwise, Bash will expand the '**\$1**' in **\$10** to the value of the first positional argument to the script.

Bash provides special variables to refer to positional parameters: **\$*** and **\$@**. Both of these variables refer to all arguments in a script, but with a slight difference. When **\$*** is used, all of the arguments are seen as a single word. However, when **\$@** is used, each argument is seen as a separate word. This is demonstrated in the following example.

```
[student@serverX bin]$ cat showargs
#!/bin/bash

for ARG in "$*"; do
    echo $ARG
done

[student@serverX bin]$ ./showargs 1 2 3
1 2 3

[student@serverX bin]$ cat showargs
#!/bin/bash

for ARG in "$@"; do
    echo $ARG
done

[student@serverX bin]$ ./showargs 1 2 3
1
2
3
```

Another value which may be useful when working with positional parameters is **\$#**, which represents the number of command-line arguments passed to a script. This value can be used to verify whether any arguments, or the correct number of arguments, are passed to a script.

```
[student@serverX bin]$ cat countargs
#!/bin/bash
echo "There are $# arguments."

[student@serverX bin]$ ./countargs
There are 0 arguments.

[student@serverX bin]$ ./countargs 1 2 3
There are 3 arguments.
```

Evaluating exit codes

Every command returns an exit status, also commonly referred to as return status or exit code. A successful command exits with an exit status of **0**. Unsuccessful commands exit with a nonzero exit status. Upon completion, a command's exit status is passed to the parent process and stored in the **?** variable. Therefore, the exit status of an executed command can be retrieved by displaying the value of **\$?**. The following examples demonstrate the execution and exit status retrieval of several common commands.

```
[student@serverX bin]$ ls /etc/hosts
/etc/hosts

[student@serverX bin]$ echo $?
0

[student@serverX bin]$ ls /etc/nofile
ls: cannot access /etc/nofile: No such file or directory

[student@serverX bin]$ echo $?
2

[student@serverX bin]$ grep localhost /etc/hosts
127.0.0.1    localhost localhost.localdomain localhost4 localhost4.localdomain4
::1        localhost localhost.localdomain localhost6 localhost6.localdomain6

[student@serverX bin]$ echo $?
0

[student@serverX bin]$ grep random /etc/hosts

[student@serverX bin]$ echo $?
1
```

Using exit codes within a script

Once executed, a script will exit when it has processed all of its contents. However, there may be times when it is desirable to exit a script midway through, such as when an error condition is encountered. This can be accomplished with the use of the **exit** command within a script. When a script encounters the **exit** command, it will exit immediately and skip the processing of the remainder of the script.

The **exit** command can be executed with an optional integer argument between **0** and **255**, which represents an exit code. An exit code value of **0** represents no error. All other nonzero values indicate an error exit code. Script authors can use different nonzero values to differentiate between different types of errors encountered. This exit code is passed back to the parent process, which stores it in the **?** variable and can be accessed with **\$?** as demonstrated in the following examples.

```
[student@serverX bin]$ cat hello
#!/bin/bash
echo "Hello, world"
exit 0

[student@serverX bin]$ ./hello
Hello, world

[student@serverX bin]$ echo $?
0
```

```
[student@serverX bin]$ cat hello
#!/bin/bash
echo "Hello, world"
exit 1

[student@serverX bin]$ ./hello
Hello, world

[student@serverX bin]$ echo $?
1
```

If the **exit** command is called without an argument, then the script will exit and pass on to the parent process the exit status of the last command executed.

Testing script inputs

To ensure that scripts are not easily derailed by unexpected conditions, it is good practice for script authors to not make assumptions regarding inputs, such as command-line arguments, user inputs, command substitutions, variable expansions, file name expansions, etc. Integrity checking can be performed by using Bash's test feature. Tests can be performed using Bash's test command syntax, [**<TESTEXPRESSION>**]. They can also be performed using Bash's newer extended test command syntax, [[**<TESTEXPRESSION>**]], which has been available since Bash version 2.02.

Like all commands, the **test** command produces an exit code upon completion, which is stored as the value **\$?** . To see the conclusion of a test, simply display the value of **\$?** immediately following the execution of the **test** command. Once again, an exit status value of **0** indicates the test succeeded, while nonzero values indicate the test failed.

Performing comparison tests

Comparison test expressions make use of *binary comparison operators*. These operators expect two objects, one on each side of the operator, and evaluate the two for equality and inequality. Bash uses a different set of operators for string and numeric comparisons, and uses the following syntax format:

```
[ <ITEM1> <BINARY COMPARISON OPERATOR> <ITEM2> ]
```

Bash's numeric comparison is limited to integer comparison. The following list of binary comparison operators is used in Bash for integer comparison.

Operator	Meaning	Example
-eq	is equal to	["\$a" -eq "\$b"]
-ne	is not equal to	["\$a" -ne "\$b"]
-gt	is greater than	["\$a" -gt "\$b"]
-ge	is greater than or equal to	["\$a" -ge "\$b"]
-lt	is less than	["\$a" -lt "\$b"]
-le	is less than or equal to	["\$a" -le "\$b"]

The following examples demonstrate the use of Bash's numeric comparison operators.


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

Bash's string comparison uses the following binary operators.

Operator	Meaning	Example
=	is equal to	["\$a" = "\$b"]
==	is equal to	["\$a" == "\$b"]
!=	is not equal to	["\$a" != "\$b"]

The following examples demonstrate the use of Bash's string comparison operators.

```
[student@serverX ~]$ [ abc = abc ]; echo $?
0
[student@serverX ~]$ [ abc == def ]; echo $?
1
[student@serverX ~]$ [ abc != def ]; echo $?
0
```

Bash also has a few unary operators available for string evaluation. Unary operators evaluate just one item using the following format.

```
[ <UNARY OPERATOR> <ITEM> ]
```

The following table shows Bash's unary operators for string evaluation.

Operator	Meaning	Example
-z	string is zero length (null)	[-z "\$a"]
-n	string is not null	[-n "\$a"]

The following examples demonstrate the use of Bash's string unary operators.

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

Testing files and directories

Bash's string and binary operators allow users to implement the good practice of not assuming the integrity of inputs to a shell script. The same caution should be utilized when scripts interact

with external entities, such as files and directories. Bash offers a large number of test operators for this purpose, as listed in the following table.

Operator	Meaning	Example
-b	file exists and is block special	[-b <FILE>]
-c	file exists and is character special	[-c <FILE>]
-d	file exists and is a directory	[-d <DIRECTORY>]
-e	file exists	[-e <FILE>]
-f	file is a regular file	[-f <FILE>]
-L	file exists and is a symbolic link	[-L <FILE>]
-r	file exists and read permission is granted	[-r <FILE>]
-s	file exists and has a size greater than zero	[-s <FILE>]
-w	file exists and write permission is granted	[-w <FILE>]
-x	file exists and execute (or search) permission is granted	[-x <FILE>]

Bash also offers a few binary comparison operators for performing file comparison. These operators are defined in the following table.

Operator	Meaning	Example
-ef	FILE1 has the same device and inode number as FILE2	[<FILE1> -ef <FILE2>]
-nt	FILE1 has newer modification date than FILE2	[<FILE1> -nt <FILE2>]
-ot	FILE1 has older modification date than FILE2	[<FILE1> -ot <FILE2>]



Important

The space characters on the inside of the brackets of the test expression, as well as those used to separate the elements within the test expression, are not there for readability, but rather are required for the proper evaluation of the expression. If any of these space characters are missing, the test will either fail or produce inaccurate or unexpected results.

Logical AND, OR operators

Situations may arise where it may be useful to test more than one condition. Bash's logical *AND* operator, `&&`, allows users to perform a compound condition test to see if both of two conditions are true. On the other hand, Bash's logical *OR* operator, `||`, allows users to test whether one of two conditions are true. The following examples demonstrate the use of Bash's logical AND and OR operators.

```
[student@serverX ~]$ [ 2 -gt 1 ] && [ 1 -gt 0 ]; echo $?
0
[student@serverX ~]$ [ 2 -gt 1 ] && [ 1 -gt 2 ]; echo $?
1
[student@serverX ~]$ [ 2 -gt 1 ] || [ 1 -gt 2 ]; echo $?
0
[student@serverX ~]$ [ 0 -gt 1 ] || [ 1 -gt 2 ]; echo $?
1
```

Using conditional structures

Simple shell scripts represent a collection of commands which are executed from beginning to end. Conditional structures allow users to incorporate decision making into shell scripts, so that certain portions of the script are executed only when certain conditions are met.

If/then statement

The simplest of the conditional structures in Bash is the `if/then` construct, which has the following syntax.

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

With this construct, if a given condition is met, one or more actions are taken. If the given condition is not met, then no action is taken. The numeric, string, and file tests previously demonstrated are frequently utilized for testing the conditions in `if/then` statements. The following code section demonstrates the use of an `if/then` statement to start the `psacct` service if it is not active.

```
systemctl is-active psacct > /dev/null 2>&1

if [ $? -ne 0 ]; then
    systemctl start psacct
fi
```

If/then/else statement

The `if/then` conditional structure can be further expanded so that different sets of actions can be taken depending on whether a condition is met. This is accomplished with the `if/then/else` conditional construct.

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

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

```
systemctl is-active psacct > /dev/null 2>&1

if [ $? -ne 0 ]; then
    systemctl start psacct
else
    systemctl stop psacct
fi
```

If/then/elif/then/else statement

Lastly, the **if/then/else** conditional structure can be further expanded to test more than one condition, executing a different set of actions when a condition is met. The construct for this is shown in the following example. In this conditional structure, Bash will test the conditions in the order presented. Upon finding a condition that is true, Bash will execute the actions associated with the condition and then skip the remainder of the conditional structure. If none of the conditions are true, then Bash will execute the actions enumerated in the **else** clause.

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

The following code section demonstrates the use of an **if/then/elif/then/else** statement to run the **mysql** client if the *mariadb* service is active, run the **psql** client if the *postgresql* service is active, or run the **sqlite3** client if both the *mariadb* and *postgresql* services are not active.

```
systemctl is-active mariadb > /dev/null 2>&1
MARIADB_ACTIVE=$?
systemctl is-active postgresql > /dev/null 2>&1
POSTGRESQL_ACTIVE=$?

if [ "$MARIADB_ACTIVE" -eq 0 ]; then
    mysql
elif [ "$POSTGRESQL_ACTIVE" -eq 0 ]; then
    psql
else
    sqlite3
fi
```

Case statement

Users can add as many **elif** clauses as they want into an **if/then/elif/then/else** statement to test as many conditions as they need. However, as more are added, the statement and its logic becomes increasingly harder to read and comprehend. For these more complex situations, Bash offers another conditional structure known as *case statements*. The **case** statement utilizes the following syntax:

```
case <VALUE> in
  <PATTERN1>
    <STATEMENT>
    ...
    <STATEMENT>
    ;;
  <PATTERN2>
    <STATEMENT>
    ...
    <STATEMENT>
    ;;
esac
```

The **case** statement attempts to match **<VALUE>** to each **<PATTERN>** in order, one by one. When a pattern matches, the code segment associated with that pattern is executed, with the **;;** syntax indicating the end of the block. All other patterns remaining in the **case** statement are then skipped and the **case** statement is exited. As many pattern/statement blocks as needed can be added.

To mimic the behavior of an **else** clause in an **if/then/elif/then/else** construct, simply use ***** as the final pattern in the **case** statement. Since this expression matches anything, it has the effect of executing a set of commands if none of the other patterns are matched.

The **case** statements are widely used in init scripts. The following code section is an example of how they are commonly used to take different actions, depending on the argument passed to the script.

```
case "$1" in
  start)
    start
    ;;
  stop)
    rm -f $lockfile
    stop
    ;;
  restart)
    restart
    ;;
  reload)
    reload
    ;;
  status)
    status
    ;;
  *)
    echo "Usage: $0 (start|stop|restart|reload|status)"
    ;;
esac
```

If the actions to be taken are the same for more than one pattern in a case statement, the patterns can be combined to share the same action block, as demonstrated in the following example. The pipe character, **|**, is used to separate the multiple patterns.

```
case "$1" in
  ...
  reload|restart)
    restart
    ;;
  ...
esac
```



References

bash(1) and **test(1)** man pages

Practice: Enhancing Bash Shell Scripts with Conditionals and Control Structures

Guided exercise

In this lab, you will create a Bash shell script to automate the process of creating Apache virtual hosts.

Resources:

Machines:

serverX

Outcomes:

A Bash script which automates the process of creating Apache virtual host configuration files and a document root directory. The script will perform the necessary checks to accommodate first-time execution of the script on a server, as well as ensure that virtual host conflicts do not occur.

Before you begin...

- Reset **serverX**.
- Log into **serverX** and become root with **sudo -i**.

Your company provides web hosting service to customers, and you have been tasked with writing a Bash shell script called `/usr/local/sbin/mkvhst` to automate the many steps involved in setting up an Apache name-based virtual host for your customers. The script will be used for virtual host creation on all servers going forward, so it needs to also be able to accommodate the one-time tasks that are executed the first time a new server is configured for name-based virtual hosting.

The script will take two arguments. The first argument will be the fully qualified domain name of the new virtual host. The second argument will be a number between 1 and 3, which represents the support tier that the customer purchased. The support tier determines the support email address, which will be set with the Apache **ServerAdmin** directive for the virtual host.

The script will create a configuration file under `/etc/httpd/conf.vhosts.d` with the name `<VIRTUALHOSTNAME>.conf` for each virtual host. It will also create a document root directory for the virtual host at `/srv/<VIRTUALHOSTNAME>/www`. Prior to creating the virtual host configuration file and document root directory, the script will check to make sure they do not already exist to ensure there will not be a conflict.

- 1. Install the `httpd` package, then enable and start `httpd`.
 - 1.1. Install the `httpd` package with `yum`.

```
[root@serverX ~]# yum install -y httpd
```

- 1.2. Enable and start `httpd`.

```
[root@serverX ~]# systemctl enable httpd
ln -s '/usr/lib/systemd/system/httpd.service' '/etc/systemd/system/multi-
user.target.wants/httpd.service'
[root@serverX ~]# systemctl start httpd
```

- 2. Begin writing your script. Store the first and second argument of the script in the **VHOSTNAME** and **TIER** variables, respectively. Set the following variables:

Variable	Value
HTTPDCONF	/etc/httpd/conf/httpd.conf
VHOSTCONFDIR	/etc/httpd/conf.vhosts.d
DEFVHOSTCONFFILE	\$VHOSTCONFDIR/00-default-vhost.conf
VHOSTCONFFILE	\$VHOSTCONFDIR/\$VHOSTNAME.conf
WWWROOT	/srv
DEFVHOSTDOCROOT	\$WWWROOT/default/www
VHOSTDOCROOT	\$WWWROOT/\$VHOSTNAME/www

- 2.1. Create the new script file with a text editor.

```
[root@serverX ~]# vim /usr/local/sbin/mkvhost
```

- 2.2. Specify the interpreter program for the script.

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

- 2.3. Set the variables for the arguments.

```
# Variables
VHOSTNAME=$1
TIER=$2
```

- 2.4. Set the other variables.

```
HTTPDCONF=/etc/httpd/conf/httpd.conf
VHOSTCONFDIR=/etc/httpd/conf.vhosts.d
DEFVHOSTCONFFILE=$VHOSTCONFDIR/00-default-vhost.conf
VHOSTCONFFILE=$VHOSTCONFDIR/$VHOSTNAME.conf
WWWROOT=/srv
DEFVHOSTDOCROOT=$WWWROOT/default/www
VHOSTDOCROOT=$WWWROOT/$VHOSTNAME/www
```

- 3. Check the argument values in the **VHOSTNAME** and **TIER** variables. If either is blank, display the message *"Usage: mkvhost VHOSTNAME TIER"* and exit with a status of **1**. If the arguments are passed correctly, then use a case statement to set a **VHOSTADMIN** variable to the proper support email address, based on the value of **\$TIER**. The case statement will use the **\$TIER** values of **1**, **2**, and **3** to set **VHOSTADMIN** to the

corresponding support email address. If any other **\$TIER** value is encountered, the case statement should display the message *"Invalid tier specified."* and exit with a status of **1**.

TIER	\$VHOSTADMIN
1	basic_support@example.com
2	business_support@example.com
3	enterprise_support@example.com

- 3.1. Create the if/then/else/fi statement. Use an OR conditional to check whether either of the arguments is blank and, if so, display the usage message and exit with a status of 1.

```
# Check arguments
if [ "$VHOSTNAME" = '' ] || [ "$TIER" = '' ]; then
    echo "Usage: $0 VHOSTNAME TIER"
    exit 1
else
```

- 3.2. Create the case statement.

```
# Set support email address
case $TIER in
    1)    VHOSTADMIN='basic_support@example.com'
          ;;
    2)    VHOSTADMIN='business_support@example.com'
          ;;
    3)    VHOSTADMIN='enterprise_support@example.com'
          ;;
    *)    echo "Invalid tier specified."
          exit 1
          ;;
esac
```

- 3.3. Close the if statement.

```
fi
```

- 4. Check to see if the **\$VHOSTCONFDIR** directory is nonexistent. If so, create the directory. Check the exit status of the directory creation and display the error message *"ERROR: Failed creating \$VHOSTCONFDIR"* if the directory creation failed.

```
# Create conf directory one time if non-existent
if [ ! -d $VHOSTCONFDIR ]; then
    mkdir $VHOSTCONFDIR

    if [ $? -ne 0 ]; then
        echo "ERROR: Failed creating $VHOSTCONFDIR."
        exit 1 # exit 1
    fi
fi
```

- 5. For Apache to be aware of the **\$VHOSTCONFDIR** directory, you must have an include statement in **\$HTTPDCONF**. Check to see if the following entry exists in the configuration

file. If not, make a backup of the file to **\$HTTPDCONF.orig** and then append the entry to the end of the configuration file. Check the exit status of the file modification and display the error message *"ERROR: Failed adding include directive"* if it failed.

```
IncludeOptional conf\vhosts\d\*\*.conf
```

```
# Add include one time if missing
grep -q '^IncludeOptional conf\vhosts\d\*\*.conf$' $HTTPDCONF

if [ $? -ne 0 ]; then
    # Backup before modifying
    cp -a $HTTPDCONF $HTTPDCONF.orig

    echo "IncludeOptional conf\vhosts\d\*\*.conf" >> $HTTPDCONF

    if [ $? -ne 0 ]; then
        echo "ERROR: Failed adding include directive."
        exit 1
    fi
fi
```

- 6. Check to see if a default virtual host already exists and, if not, create it.
 - 6.1. Verify if the default virtual host configuration file already exists and, if not, create and populate it with the following statement:

```
cat <<DEFCONFEOF > $DEFVHOSTCONFFILE
<VirtualHost _default_:80>
    DocumentRoot $DEFVHOSTDOCROOT
    CustomLog "logs/default-vhost.log" combined
</VirtualHost>

<Directory $DEFVHOSTDOCROOT>
    Require all granted
</Directory>
DEFCONFEOF
```

```
# Check for default virtual host
if [ ! -f $DEFVHOSTCONFFILE ]; then
    cat <<DEFCONFEOF > $DEFVHOSTCONFFILE
<VirtualHost _default_:80>
    DocumentRoot $DEFVHOSTDOCROOT
    CustomLog "logs/default-vhost.log" combined
</VirtualHost>

<Directory $DEFVHOSTDOCROOT>
    Require all granted
</Directory>
DEFCONFEOF
fi
```

- 6.2. Verify if the default virtual host document root directory already exists and, if not, create it.

```

if [ ! -d $DEFVHOSTDOCROOT ]; then
    mkdir -p $DEFVHOSTDOCROOT
    restorecon -Rv /srv/
fi

```

- 7. Check to see if the virtual host's configuration file already exists and, if so, display the error message *"ERROR: \$VHOSTCONFFILE already exists."* and exit with a status of **1**. Check to see if the virtual host's document root directory already exists and, if so, display the error message *"ERROR: \$VHOSTDOCROOT already exists."* and exit with a status of **1**. If no errors are encountered with the previous two checks, continue with the creation of the virtual host configuration file, **\$VHOSTCONFFILE**, and document root directory, **\$VHOSTDOCROOT**. Populate the virtual host configuration file with the following here statement:

```

cat <<CONFEOF > $VHOSTCONFFILE
<VirtualHost *:80>
    ServerName $VHOSTNAME
    ServerAdmin $VHOSTADMIN
    DocumentRoot $VHOSTDOCROOT
    ErrorLog "logs/${VHOSTNAME}_error_log"
    CustomLog "logs/${VHOSTNAME}_access_log" common
</VirtualHost>

<Directory $VHOSTDOCROOT>
    Require all granted
</Directory>
CONFEOF

```

```

# Check for virtual host conflict
if [ -f $VHOSTCONFFILE ]; then
    echo "ERROR: $VHOSTCONFFILE already exists."
    exit 1
elif [ -d $VHOSTDOCROOT ]; then
    echo "ERROR: $VHOSTDOCROOT already exists."
    exit 1
else
    cat <<CONFEOF > $VHOSTCONFFILE
<Directory $VHOSTDOCROOT>
    Require all granted
    AllowOverride None
</Directory>

<VirtualHost *:80>
    DocumentRoot $VHOSTDOCROOT
    ServerName $VHOSTNAME
    ServerAdmin $VHOSTADMIN
    ErrorLog "logs/${VHOSTNAME}_error_log"
    CustomLog "logs/${VHOSTNAME}_access_log" common
</VirtualHost>
CONFEOF

    mkdir -p $VHOSTDOCROOT
    restorecon -Rv $WWWROOT
fi

```

- 8. Verify the syntax of the configuration file with the command **apachectl configtest**. If there are no errors, then reload **httpd.service**. Otherwise, display the error message *"ERROR: Configuration error."* and exit with a status of **1**.

```
# Check config and reload
apachectl configtest &> /dev/null

if [ $? -eq 0 ]; then
    systemctl reload httpd &> /dev/null
else
    echo "ERROR: Config error."
    exit 1
fi
```

- 9. Save and execute the script.
 - 9.1. Make the script executable.

```
[root@serverX ~]# chmod u+x /usr/local/sbin/mkvhost
```

- 9.2. Execute the script.

```
[root@serverX ~]# /usr/local/sbin/mkvhost wwwX.example.com 3
```

Lab: Bash Conditionals and Control Structures

Performance checklist

In this lab, you will enhance and add new features to a Bash script that automates the creation of user accounts.

Resources:	
Files:	<ul style="list-style-type: none"> • /usr/local/sbin/mkaccounts • /tmp/support/newusers
Machines:	serverX

Outcomes:

A shell script, `/usr/local/sbin/mkaccounts`, which creates new user accounts based on information provided in a data file. The shell script will support two different options for an argument to offer a verbose mode or to display a usage message. Prior to the creation of new accounts, the script will check for and report conflicts with existing accounts.

Before you begin...

- Reset **serverX**.
- Log into **serverX** and become the **root** user with **sudo -i**.
- Run **lab bashconditionals setup**.

Your company provides hosting service to customers, and currently uses a Bash shell script called `/usr/local/sbin/mkaccounts` to automate the daily task of creating accounts for new customers by processing a colon-separated data file located at `/tmp/support/newusers`.

You have been tasked with extending the script to add the following new features:

- The script should accept a command-line argument so that it can either be run in verbose mode or generate a usage message.
- When in verbose mode, the script should generate a message to indicate the creation of each account.
- Prior to the creation of an account, the script should check existing accounts and report if a conflict or duplication occurs.

Once you have completed the script, run **lab bashconditionals grade** to verify its functionality.

1. Begin by extending the script located at `/usr/local/sbin/mkaccounts` by storing the first argument to the **OPTION** variable.
2. After the variable declaration at the top of the script, add a case statement that expects three patterns for the value of the **OPTION** variable and takes the following corresponding actions.

Pattern	Action
Blank value	No action.
-v	Set variable VERBOSE to value of 'y'.
-h	Display usage message " <i>Usage: mkaccounts [-h -v]</i> " and exit.
All other patterns	Display usage message " <i>Usage: mkaccounts [-h -v]</i> " and exit with status 1 .

3. Within the existing for loop, prior to creating the new user account, check whether an account of the same name already exists. If an account already exists, set the value of the **ACCTEXIST** variable to 'y'. Store the value of the *GECOS* field of the existing account in the variable **ACCTEXISTNAME**.
4. Within the existing for loop, use the **ACCTEXIST** and **ACCTEXISTNAME** variables to perform the following evaluations and corresponding actions.

Condition	Action
Account exists and the full name of the existing account matches that of a new customer.	Display message " <i>Skipping <ACCOUNTNAME>. Duplicate found.</i> "
Account exists and the full name of the existing account does not match that of a new customer.	Display message " <i>Skipping <ACCOUNTNAME>. Conflict found.</i> "
Account does not already exist.	Create the account. If the value of the variable VERBOSE is set to 'y', display message " <i>Added <ACCOUNTNAME>.</i> "

5. Save and execute the script. Verify that the script performed as intended with each of the expected options. You may need to run **lab bashconditionals clean** to purge out accounts that might have been created before rerunning your script.
6. Verify your script by running **lab bashconditionals grade**.

Solution

In this lab, you will enhance and add new features to a Bash script that automates the creation of user accounts.

Resources:	
Files:	<ul style="list-style-type: none"> • /usr/local/sbin/mkaccounts • /tmp/support/newusers
Machines:	serverX

Outcomes:

A shell script, `/usr/local/sbin/mkaccounts`, which creates new user accounts based on information provided in a data file. The shell script will support two different options for an argument to offer a verbose mode or to display a usage message. Prior to the creation of new accounts, the script will check for and report conflicts with existing accounts.

Before you begin...

- Reset **serverX**.
- Log into **serverX** and become the **root** user with **sudo -i**.
- Run **lab bashconditionals setup**.

Your company provides hosting service to customers, and currently uses a Bash shell script called `/usr/local/sbin/mkaccounts` to automate the daily task of creating accounts for new customers by processing a colon-separated data file located at `/tmp/support/newusers`.

You have been tasked with extending the script to add the following new features:

- The script should accept a command-line argument so that it can either be run in verbose mode or generate a usage message.
- When in verbose mode, the script should generate a message to indicate the creation of each account.
- Prior to the creation of an account, the script should check existing accounts and report if a conflict or duplication occurs.

Once you have completed the script, run **lab bashconditionals grade** to verify its functionality.

1. Begin by extending the script located at `/usr/local/sbin/mkaccounts` by storing the first argument to the **OPTION** variable.
 - 1.1. Back up the script prior to making edits.

```
[root@serverX ~]# cp -a /usr/local/sbin/mkaccounts /usr/local/sbin/mkaccounts.orig
```

- 1.2. After the existing line declaring the **NEWUSERSFILE** variable, add a new variable, **OPTION**, to store the first argument.

```
OPTION=$1
```

2. After the variable declaration at the top of the script, add a case statement that expects three patterns for the value of the **OPTION** variable and takes the following corresponding actions.

Pattern	Action
Blank value	No action.
-v	Set variable VERBOSE to value of 'y'.
-h	Display usage message " <i>Usage: mkaccounts [-h -v]</i> " and exit.
All other patterns	Display usage message " <i>Usage: mkaccounts [-h -v]</i> " and exit with status 1 .

```
case $OPTION in
  '')
    ;;
  -v)
    VERBOSE=y
    ;;
  -h)
    echo "Usage: $0 [-h|-v]"
    echo
    exit
    ;;
  *)
    echo "Usage: $0 [-h|-v]"
    echo
    exit 1
    ;;
esac
```

3. Within the existing for loop, prior to creating the new user account, check whether an account of the same name already exists. If an account already exists, set the value of the **ACCTEXIST** variable to 'y'. Store the value of the **GECOS** field of the existing account in the variable **ACCTEXISTNAME**.

```
# Test for dups and conflicts
ACCTEXIST=''
ACCTEXISTNAME=''

id $ACCTNAME &> /dev/null

if [ $? -eq 0 ]; then
  ACCTEXIST=y
  ACCTEXISTNAME="$(grep ^$ACCTNAME: /etc/passwd | cut -f5 -d:)"
fi
```

4. Within the existing for loop, use the **ACCTEXIST** and **ACCTEXISTNAME** variables to perform the following evaluations and corresponding actions.

Condition	Action
Account exists and the full name of the existing account matches that of a new customer.	Display message " <i>Skipping <ACCOUNTNAME>. Duplicate found.</i> "
Account exists and the full name of the existing account does not match that of a new customer.	Display message " <i>Skipping <ACCOUNTNAME>. Conflict found.</i> "
Account does not already exist.	Create the account. If the value of the variable VERBOSE is set to 'y', display message " <i>Added <ACCOUNTNAME>.</i> "

```
# Test for dups and conflicts
if [ "$ACCTEXIST" = 'y' ] && [ "$ACCTEXISTNAME" = "$FIRSTNAME
$LASTNAME" ]; then
    echo "Skipping $ACCTNAME. Duplicate found."
elif [ "$ACCTEXIST" = 'y' ]; then
    echo "Skipping $ACCTNAME. Conflict found."
else
    useradd $ACCTNAME -c "$FIRSTNAME $LASTNAME"

    if [ "$VERBOSE" = 'y' ]; then
        echo "Added $ACCTNAME."
    fi
fi
```

5. Save and execute the script. Verify that the script performed as intended with each of the expected options. You may need to run **lab bashconditionals clean** to purge out accounts that might have been created before rerunning your script.

- 5.1. Make the script executable.

```
[root@serverX ~]# chmod u+x /usr/local/sbin/mkaccounts
```

- 5.2. Execute the script with an invalid argument.

```
[root@serverX ~]# /usr/local/sbin/mkaccounts test
Usage: ./mkaccounts [-h] -v]

[root@serverX sbin]# echo $?
1
```

- 5.3. Execute the script with the **-h** argument.

```
[root@serverX ~]# /usr/local/sbin/mkaccounts -h
Usage: ./mkaccounts [-h] -v]

[root@serverX sbin]# echo $?
0
```

- 5.4. Execute the script with the **-v** argument.

```
[root@serverX ~]# /usr/local/sbin/mkaccounts -v
Added bwerts.
Added hballa.
...
Added crundle.
Added tgrana.
Skipping bwerts. Conflict found.
Skipping hballa. Duplicate found.
"Tier 1","24","46%"
"Tier 2","15","28%"
"Tier 3","13","25%"
```

- 5.5. Execute the script with no arguments. Since accounts have already been created, they will first need to be purged with **lab bashconditionals clean**.

```
[root@serverX ~]# lab bashconditionals clean
[root@serverX ~]# /usr/local/sbin/mkaccounts
Skipping bwerts. Conflict found.
Skipping hballa. Duplicate found.
"Tier 1","24","46%"
"Tier 2","15","28%"
"Tier 3","13","25%"
```

6. Verify your script by running **lab bashconditionals grade**.

Summary

Enhancing Bash Shell Scripts with Conditionals and Control Structures

In this section, students learned how to:

- Extend script functionality with the use of positional parameters.
- Evaluate exit status to verify success or failure of program execution.
- Perform tests on script inputs, files, and directories with the **test** command.
- Utilize **if/then** and **case** statements to perform different actions for different conditions.



CHAPTER 13

CONFIGURING THE SHELL ENVIRONMENT

Overview	
Goal	To customize Bash startup and use environment variables, Bash aliases, and Bash functions.
Objectives	<ul style="list-style-type: none">• Use bash startup scripts to define environment variables, aliases, and functions
Sections	<ul style="list-style-type: none">• Changing the Shell Environment (and Practice)
Lab	<ul style="list-style-type: none">• Configuring the Shell Environment

Changing the Shell Environment

Objectives

After completing this section, students should be able to describe environment variables, configure them in **bash** start-up scripts, define **bash** aliases, and use functions.

Environment variables

The shell and scripts use variables to store data; some variables can be passed to sub-processes along with their content. These special variables are called *environment variables*.

Applications and sessions use these variables to determine their behavior. Some of them are likely familiar to administrators, such as **PATH**, **USER**, and **HOSTNAME**, among others. What makes variables environment variables is that they have been *exported* in the shell. A variable that is flagged as an export will be passed, with its value, to any sub-process spawned from the shell. Users can use the **env** command to view all environment variables that are defined in their shell.

Any variable defined in the shell can be an environment variable. The key to making a variable become an environment variable is flagging it for export using the **export** command.

In the following example, a variable, **MYVAR**, will be set. A sub-shell is spawned, and the **MYVAR** variable does not exist in the sub-shell.

```
[student@demo ~]$ MYVAR="some value"
[student@demo ~]$ echo $MYVAR
some value
[student@demo ~]$ bash
[student@demo ~]$ echo $MYVAR
[student@demo ~]$ exit
```

In a similar example, the **export** command will be used to tag the **MYVAR** variable as an environment variable, which will be passed to a sub-shell.

```
[student@demo ~]$ MYVAR="some value"
[student@demo ~]$ export MYVAR
[student@demo ~]$ echo $MYVAR
some value
[student@demo ~]$ bash
[student@demo ~]$ echo $MYVAR
some value
[student@demo ~]$ exit
```

bash start-up scripts

One place where environment variables are used is in initializing the **bash** environment upon user log in. When a user logs in, several shell scripts are executed to initialize their environment, starting with the **/etc/profile**, followed by a profile in the user's home directory, typically **~/.bash_profile**.



Note

A **bash** shell will look for one of three files in a user's home directory: **.bash_profile**, **.bash_login**, or **.profile**. The shell will look for the files in the given order, and will execute the first file it locates.

Because these profiles have additional scripting in them, which calls other shell scripts, the **bash** login scripting will typically be the following:

```
/etc/profile
    \__ /etc/profile.d/*.sh

~/.bash_profile
    \__ ~/.bashrc
        \__ /etc/bashrc
```

Generally, there are two types of login scripts, profiles and "RCs". Profiles are for setting and exporting of environment variables, as well as running commands that should only be run upon login. RCs, such as **/etc/bashrc**, are for running commands, setting aliases, defining functions, and other settings that cannot be exported to sub-shells. Usually, profiles are only executed in a login shell, whereas RCs are executed every time a shell is created, login or non-login.

The layout of the file call is such that a user can override the default settings provided by the system wide scripts. Many of the configuration files provided by Red Hat will contain a comment indicating where user-specific changes should be added.

Using **alias**

alias is a way administrators or users can define their own command to the system or override the use of existing system commands. Aliases are parsed and substituted prior to the shell checking **PATH**. **alias** can also be used to display all existing aliases defined in the shell.

```
[student@demo ~]$ alias
alias egrep='egrep --color=auto'
alias fgrep='fgrep --color=auto'
alias grep='grep --color=auto'
alias l.='ls -d .* --color=auto'
alias ll='ls -l --color=auto'
alias ls='ls --color=auto'
alias vi='vim'
alias which='alias | /usr/bin/which --tty-only --read-alias --show-dot --show-tilde'
```

All the default aliases defined in a user environment are in the previously shown output of **alias**. The definition of the **ll** alias means that when the user types **ll** as a command, the shell will expand the alias and execute **ls -l --color=auto**. In this way, a new command **ll** has been added to the shell. In another example, the alias for **grep**, the alias overrides the default invocation of an existing command on the system. When a user enters the **grep** command, the shell will expand the alias and substitute the **grep --color=auto** command. Due to the alias, all calls to **grep** are overridden to become calls to **grep** with the **--color** option passed automatically.

Use the **alias** command to set an alias. The defined alias will exist for the duration of the current shell only.

```
alias mycomm="<command to execute>"
```

```
[student@demo ~]$ alias usercmd='echo "Hurrah!"; ls -l'
[student@demo ~]$ usercmd
Hurrah!
total 0
-rw-rw-r--. 1 student student 0 Jun  9 13:21 file1
-rw-rw-r--. 1 student student 0 Jun  9 13:21 file2
-rw-rw-r--. 1 student student 0 Jun  9 13:21 file3
```

To make the alias persistent, the user would need to add the command to the bottom of their `~/.bashrc`.

```
[student@demo ~]$ vi ~/.bashrc
...
# User specific aliases and functions
alias usercmd='echo "Hurrah!"; ls -l'
```

After the **alias** is added to the `~/.bashrc`, it will be available in every shell created.

To remove an alias from the environment, use the **unalias** command.

Using functions

When used in shell scripts, functions are a way of isolating a segment of code so that it can be called repeatedly without having to write the entire segment again. Additionally, if the code requires an update, the function's content can be updated, and everywhere the function is referenced, the updated code is now executed.

An example of defining and using a function within a shell script follows, taken from `/etc/profile`. The **pathmunge** function takes two arguments; the first is a directory, the second (optional) is the word "after". Based on whether "after" is passed as `$2`, the directory will be added to the **PATH** environment variable at the front or end of the existing list of directories. Later, the function is invoked several times to build the **PATH** for root or regular users. Notice that for root, all the directories are prepended to **PATH**, where regular users have their **PATH** built by appending.

```
pathmunge () {
  if [ "$2" = "after" ] ; then
    PATH=$PATH:$1
  else
    PATH=$1:$PATH
  fi
}
...
if [ "$EUID" = "0" ]; then
  pathmunge /sbin
  pathmunge /usr/sbin
  pathmunge /usr/local/sbin
else
  pathmunge /usr/local/sbin after
  pathmunge /usr/sbin after
  pathmunge /sbin after
fi
```

Functions can also be set in the **bash** shell environment. When set in the environment, they can be executed as commands on the command line, similar to aliases. Unlike aliases, they can take arguments, be much more sophisticated in their actions, and provide a return code. Functions can be defined in the current shell by typing them into the command line, but more realistically, they should be set in a user's `~/ .bashrc` or the global `/etc/bashrc`.

There are many functions set by default in the user environment, which can be viewed with the **set** command. **set** will display all functions and variables in the current shell environment. To remove a function from the environment, a user or administrator may use the **unset** command, passing the function or variable name to remove.



References

bash(1), **env(1)**, and **builtins(1)** man pages

Practice: Working with Login and Non-Login Shells

Guided exercise

In this lab, you will apply changes to the scripts used to establish the initial shell environment for log in and non-login shells.

Resources:	
Files:	<ul style="list-style-type: none"> • /etc/profile • /etc/bashrc • ~/.bash_profile • ~/.bashrc
Machines:	serverX

Outcomes:

Shells will now have additional environment variables, functions, and aliases defined.

Before you begin...

Reset your **serverX** system. After the reset completes, log in as the **student** user.

- 1. Change the student user's **PS1** environment variable to `[\u@\h \t \w]$.`
 - 1.1. Since this is an environment variable, and it should affect only student, edit `~/.bash_profile`.

```
[student@serverX ~]$ vi ~/.bash_profile
```

- 1.2. Add an entry for the **PS1** variable, setting it to `[\u@\h \t \w]$.`

```
...
# User specific environment and startup programs

PATH=$PATH:$HOME/.local/bin:$HOME/bin
PS1='[\u@\h \t \w]$ '

export PATH PS1
```

- 2. Set aliases for the student user so that when the **rm**, **cp**, or **mv** commands are used, they are automatically called with the **-i** option.
 - 2.1. Edit the student's `~/.bashrc` configuration file.

```
[student@serverX ~]$ vi ~/.bashrc
```

- 2.2. Add aliases for the commands into the file.

```
...
# User specific aliases and functions

alias rm='rm -i'
alias cp='cp -i'
alias mv='mv -i'
```

- 3. You have decided that users on serverX need to have an application to help them divine the future. Add a function called **8ball**, available to all users, that selects a random number between zero and three. Based on the random value, the function should present a message to help them make a decision about the future.

- 3.1. Edit the **/etc/bashrc** file.

```
[student@serverX ~]$ sudo vi /etc/bashrc
```

- 3.2. At the bottom of the file, add the **8ball** function.

```
...
8ball () {
    echo "Shaking ..."
    echo
    sleep 3
    value="$[ $RANDOM % 3 ]"
    case $value in
        0) echo "All signs point to yes." ;;
        1) echo "The answer is no." ;;
        2) echo "Ask again later." ;;
        3) echo "Outlook hazy." ;;
    esac
    echo
}
}
```

- 4. Use an **ssh** session to create a login shell to verify that the environment changes are effective.

```
[student@serverX ~]$ ssh student@localhost
student@localhost's password: student
[student@serverX 00:06:57 ~]$ touch file1
[student@serverX 00:06:57 ~]$ rm file1
rm: remove regular empty file 'file1'? y
[student@serverX 00:06:57 ~]$ 8ball
The answer is no.
```

Lab: Configuring the Shell Environment

Performance checklist

In this lab, you will make configuration changes to the shell environment for individual users, and all users, on the machine.

Your international coworkers have complained that when they log into the console of serverX, the language on the machine is set incorrectly. You will configure the machine such that when a user logs in with the terminal type of **linux**, the **LANG** variable will be set to **en_US**.

The student user requires a command be added to their environment. **diskcheck** will run **iostat -d** and **df -hP --type xfs**.

The student user needs an environment variable **JAVA_HOME** set to **/usr/lib/jvm**.

Resources:	
Files:	<ul style="list-style-type: none"> • /etc/profile • ~/.bash_profile • ~/.bashrc
Machines:	serverX

Outcomes:

All users will have their **LANG** variable set to **en_US** when using a **linux** terminal type. The student user will have a function, named **diskcheck**, defined in the environment. The student will have an environment variable, **JAVA_HOME**, defined in the environment.

Before you begin...

Reset your **serverX** system.

1. Detect the terminal type of a shell. If it is **linux**, make sure the language is set to **en_US**.
2. Add a function to the student user's environment. The function is called **diskcheck**, and when called, will display the output of the **iostat -d** and **df -hP --type xfs** commands.
3. Set an environment variable, **JAVA_HOME**, to **/usr/lib/jvm** for the student user.

Solution

In this lab, you will make configuration changes to the shell environment for individual users, and all users, on the machine.

Your international coworkers have complained that when they log into the console of serverX, the language on the machine is set incorrectly. You will configure the machine such that when a user logs in with the terminal type of **linux**, the **LANG** variable will be set to **en_US**.

The student user requires a command be added to their environment. **diskcheck** will run **iostat -d** and **df -hP --type xfs**.

The student user needs an environment variable **JAVA_HOME** set to **/usr/lib/jvm**.

Resources:	
Files:	<ul style="list-style-type: none"> • /etc/profile • ~/.bash_profile • ~/.bashrc
Machines:	serverX

Outcomes:

All users will have their **LANG** variable set to **en_US** when using a **linux** terminal type. The student user will have a function, named **diskcheck**, defined in the environment. The student will have an environment variable, **JAVA_HOME**, defined in the environment.

Before you begin...

Reset your **serverX** system.

1. Detect the terminal type of a shell. If it is **linux**, make sure the language is set to **en_US**.

Edit the **/etc/profile** file. At the bottom of the file, add a bit of script which will check the **TERM** variable to see if it is set to **linux**. If it is, assign the **LANG** variable to **en_US** and **export** the **LANG** setting.

```
if [ "$TERM" == "linux" ]
then
  LANG=en_US
  export LANG
fi
```

To verify the setting works, switch to a virtual console and log in as a user. Check the value of the **LANG** variable.

```
[student@serverX ~]$ echo $LANG
en_US
```

2. Add a function to the student user's environment. The function is called **diskcheck**, and when called, will display the output of the **iostat -d** and **df -hP --type xfs** commands.

Edit the student user's `~/ .bashrc` and add the function definition. Add the new function at the end of the file.

```
# User specific aliases and functions
diskcheck() {
    iostat -d
    echo
    df -hP --type xfs
}
```

Source the `~/ .bashrc` and verify the function.

```
[student@serverX ~]$ . ~/.bashrc
[student@serverX ~]$ diskcheck
```

3. Set an environment variable, `JAVA_HOME`, to `/usr/lib/jvm` for the student user.

Edit the student user's `.bash_profile`, and at the bottom of the file, add an entry for `JAVA_HOME`. Set `JAVA_HOME` to `/usr/lib/jvm`. Use the `export` command to tag the variable to be available for all sub-shells.

```
JAVA_HOME=/usr/lib/jvm
export JAVA_HOME
```

Source the `.bash_profile` to read the changes into the environment.

```
[student@serverX ~]$ . ~/.bash_profile
```

Ensure the variable is available in sub-shells.

```
[student@serverX ~]$ bash
[student@serverX ~]$ echo $JAVA_HOME
/usr/lib/jvm
```

Summary

Changing the Shell Environment

In this section, students learned how to configure **bash** start up scripts to define environment variables, aliases, and functions.





CHAPTER 14

LINUX CONTAINERS AND DOCKER

Overview	
Goal	To manage Linux containers with Docker in Red Hat Enterprise Linux 7.
Objectives	<ul style="list-style-type: none">• Explain what Linux containers are and Docker's role in managing them.• Manage Linux containers using Docker.
Sections	<ul style="list-style-type: none">• Introduction to Linux Containers• Using Docker

Introduction to Linux Containers

Objectives

After completing this section, students should be able to explain what Linux containers are and the role Docker plays in managing them.

Linux containers

A *container* is a lightweight application isolation mechanism that allows the kernel to run groups of processes in their own isolated user spaces, separate from the host system. The container has its own process list, network stack, file systems, and other resources, but shares the kernel with the host and the other containers running on the system. Linux containers are implemented through a combination of three kernel features: *namespaces* for isolation, *control groups* for resource control, and *SELinux* for security.

A tool called **Docker** is used to create, control, and manage containers. **Docker** adds an API, an image format, and a delivery and sharing model to Linux containers. A *Docker image* contains an application and all its dependencies. When a container is started, a read-write layer for that container is combined with the read-only image using LVM thin provisioning. Docker images are portable and can be saved and exported to other systems and users.



Important

At the time of writing (June 2014), Linux containers and Docker in Red Hat Enterprise Linux 7 were rapidly evolving technologies. Some details of this chapter may change in the weeks and months to come. Containers are fully supported, but due to the rapid ongoing development of these tools, Red Hat requests that customers do not deploy Docker in a mission-critical environment at this time.

For up-to-date documentation on containers in RHEL 7, and future updates on Red Hat's support policy, please see the references posted on access.redhat.com listed at the end of this section.

This content is currently not an RHCE exam competency.

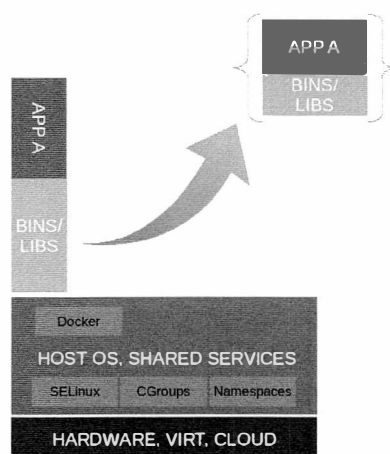


Figure 14.1: Docker, containers, and Docker images

Container technology

Namespaces

The kernel provides container isolation through namespaces, which create a new environment with a unique view of a subset of the resources on the system. Some have called namespaces "**chroot** on steroids," since instead of just providing processes with a different root file system, processes also have a different view of the process table, networking stack, and so on.

There are five different types of namespaces currently in use by containers:

Namespaces Used by Linux Containers

Namespace	Purpose
Mount	Isolates the file systems seen by the container. Containers have a different / than each other, and also different than the host system.
PID	Each container has its own process ID (PID) table. Processes inside a container cannot see outside processes, and PID 1 is different in each container. All processes in the container are visible from the host, but will have different PID numbers.
Network	Each container has its own network interfaces, routing table, and firewall rules, which are not directly visible from the host's default network namespace or each other. These can be connected to the host's network infrastructure and the outside world.
IPC	Isolates the interprocess communication (IPC) resources such as System V shared memory and POSIX message queues. Two containers cannot interact with each other's shared memory spaces.
UTS	A container can have a different host name and domain name than other containers and the host system.

Control groups

Control groups (cgroups) are used by the kernel to manage system resources. Cgroups allow fair (or unfair) allocation of CPU time, memory, and I/O bandwidth among processes and groups of processes. Containers use cgroups to manage resource consumption, so that a container will get a certain share of system resources but not steal all system resources.

Management of cgroups is beyond the scope of this course. System administrators familiar with cgroups should be aware that RHEL 7 now uses **systemd** scope and slice units to more easily manage cgroups.

SELinux and containers

Not everything is namespaced, and containers are less secure than virtual machines. In order to protect the host and other containers from a compromised container, SELinux is used.

With SELinux enforcing, container processes can only write to container files. Container processes run as the type **svirt_lxc_net_t**, and image files are labeled with the type **svirt_sandbox_file_t**.

To keep two containers from accessing each other, even though they have the same SELinux type, Multi-Category Security (MCS) types are used. Each container runs as **svirt_lxc_net_t**, but is assigned a random and unique *category* when it is started by Docker. If a process has the

right SELinux type but its category does not match, access is denied. A reference at the end of this section provides a high-level overview of SELinux MCS type checking.



Warning

Always run the host system with SELinux in enforcing mode.

If the host system being used to control containers is being run with SELinux disabled or set to permissive mode, it is possible for a hostile process running in a container to break out and compromise the host environment and all other containers on the system.

Docker

Docker provides the user interface, API, image format, and other tools used to manage Linux containers in RHEL 7. The *docker* package is provided by a special repository, RHEL 7 Extras (**rhel-7-server-extras-rpms**), which is updated under more aggressive engineering criteria than the core distribution.

Docker runs a daemon started by the **systemd** unit **docker.service** to manage containers. The user interacts with this daemon through the **docker** command. Images are stored in a local *index* kept in the **/var/lib/docker** directory, but are loaded and exported with the **docker** command.

Docker images

A Docker *image* is a static snapshot of a container's configuration, which is used to launch a container. The image is a read-only layer that is never modified. Instead, Docker adds a read-write overlay to which all changes are made. Changes are saved by creating a new image. A single image can be used to generate many containers that are very slightly different, but only need enough disk space to store a very small amount of differences.

A *platform image* is an image with no parent. This is a baseline image that defines the core runtime environment, packages, and utilities needed to run an application. Red Hat provides an extremely minimal platform image for RHEL 7 that has the minimum required packages to support a container, **bash**, and to use **yum** to download and install additional packages into a container.

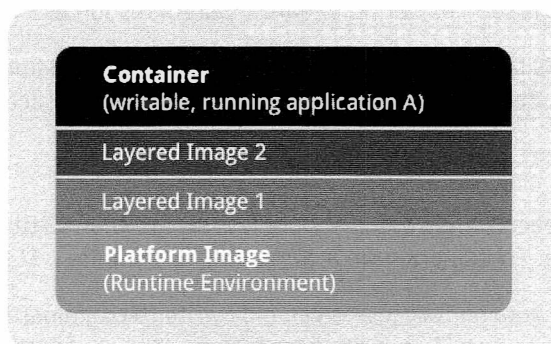


Figure 14.2: Image layering using Docker

The Docker registry

Platform images can be provided as a **tar** archive and loaded manually. More frequently, they are *pulled* from a *Docker registry*. This repository may be available to the public. Red Hat provides

a read-only registry at `registry.access.redhat.com`, from which the RHEL 7 base platform image can be pulled. Docker maintains a public registry from which open source community-provided images can be pulled.

It is also possible to use the `docker-registry` package to set up a simple private repository. Images can be pushed to this private repository and pulled by other machines in the network.

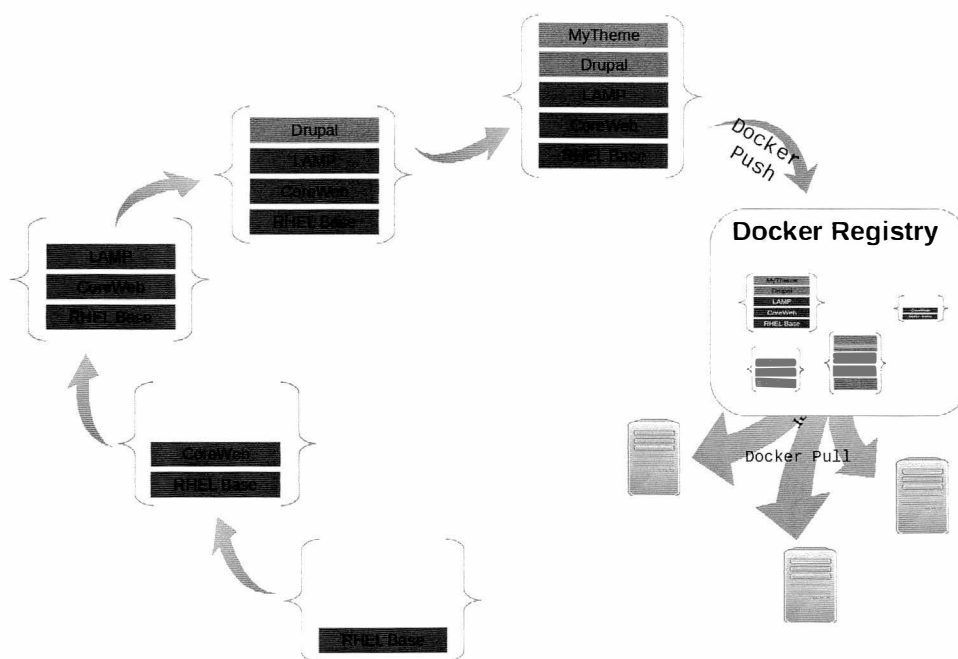


Figure 14.3: Docker images and a Docker registry

Containers and virtualization

Containers and virtualization are two technologies that provide complementary ways of dividing up system resources. Containers can be run on virtual machines and in cloud computing environments, combining both technologies. If virtualization “vertically” abstracts hardware, containers can be said to “horizontally” segment operating systems.

There are use cases more suited for one technology or the other. Some advantages of Docker containers:

- Containers are lightweight in resource use, so more containers can be run on the same hardware.
- Containers can be created and destroyed more quickly than virtual machines.
- Unlike virtual machines, containers do not need to support an entire operating system; only a core runtime is needed for the application. This allows rapid application deployment.
- Docker images have a version control stream, so successive versions of an image can be tracked and even reverted. Components reuse components from other layers, making container images very lightweight.
- Docker images are easily transferable to other machines which have Docker available.

By comparison, some advantages of virtualization:

- Virtual machines run their own kernel and full operating system, which allows stronger isolation between the host hypervisor and the virtual machine.
- Virtual machines can easily run operating systems and kernels that are completely different than the hypervisor host's operating system.
- Virtual machines can be live-migrated from one hypervisor node to another while running; containers must be stopped before being moved from one machine to another.

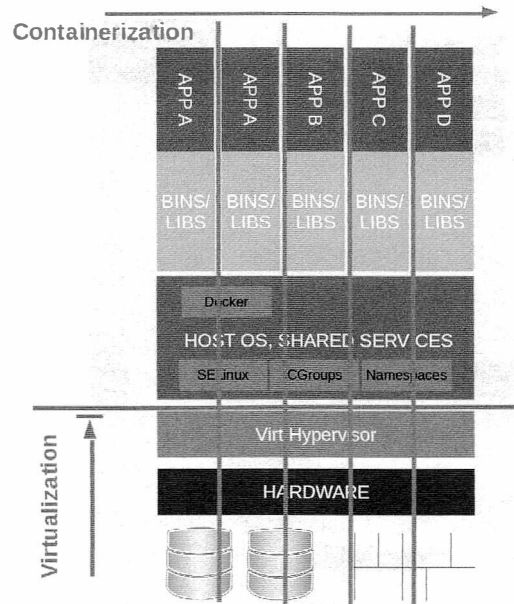


Figure 14.4: Containers and virtualization

Red Hat Enterprise Linux Atomic Host

Red Hat is working with Project Atomic (<http://www.projectatomic.io>) to produce a variant of Red Hat Enterprise Linux 7 (*Red Hat Enterprise Linux Atomic Host*) that is optimized to run containers in a small footprint, and which can be updated atomically for easier maintainability. Running containers on a Project Atomic system is somewhat analogous to running virtual machines on a thin hypervisor host. Only the components needed to support the containers would be provided by the host environment.



References

FAQ: What is Docker and how is it integrated in RHEL7?
<https://access.redhat.com/site/solutions/906143>

How is Docker supported in RHEL7?
<https://access.redhat.com/site/solutions/907283>

Get Started with Docker Containers in RHEL7
<https://access.redhat.com/site/articles/881893>

Red Hat Enterprise Linux 7 - Docker Overview
<https://access.redhat.com/site/videos/911053>

Docker
<http://www.docker.com>

Project Atomic
<http://www.projectatomic.io>

More information is available in the "Linux Containers with Docker Format" section of the *Red Hat Enterprise Linux 7.0 Release Notes* and in the *Red Hat Enterprise Linux 7 Resource Management and Linux Containers Guide*, which can both be found at <https://access.redhat.com/docs>

SELinux plays an important role in Linux container security. The following reference is a useful guide to illustrate SELinux policy concepts, including MCS and MLS, in a light-hearted way:

The SELinux Coloring Book
<http://blog.linuxgrrl.com/2014/04/16/the-selinux-coloring-book/>

Using Docker

Objectives

After completing this section, students should be able to use Docker to manage Linux containers.

Installing and starting Docker

To download and use the Docker package that comes with Red Hat Enterprise Linux 7, the RHEL 7 system needs to be registered and properly subscribed to RHEL 7 using **subscription-manager**. Once that is done, access to the RHEL 7 Optional and RHEL 7 Extras channels needs to be enabled and the *docker* package installed.

```
[root@demo ~]# subscription-manager repos --enable=rhel-7-server-extras-rpms
[root@demo ~]# subscription-manager repos --enable=rhel-7-server-optional-rpms
[root@demo ~]# yum install docker
```



Warning

At the time of writing, a bug in Docker required that the **firewalld** service be turned off in order to run Docker, https://bugzilla.redhat.com/show_bug.cgi?id=1098281. It is expected that this issue will be resolved in the near future, but in the meantime, the service must be disabled.

```
[root@demo ~]# systemctl stop firewalld
[root@demo ~]# systemctl disable firewalld
```

Once Docker is installed, the **docker** service must be started and enabled.

```
[root@demo ~]# systemctl start docker; systemctl enable docker
```

Pulling or loading platform images

Once Docker is installed, the next step is to get base platform images into the machine's local index. One way to do this is to pull the images from a remote Docker registry. This can be done with the command **docker pull repository/image-name**. If no repository is specified, the community Docker hub repository is assumed. Red Hat has an official base image named **redhat/rhel7** available at registry.access.redhat.com.

```
[root@demo ~]# docker pull registry.access.redhat.com/redhat/rhel7
Pulling repository registry.access.redhat.com/redhat/rhel7
e1f5733f050b: Download complete
```

Another way to get base images is to get them in the form of an exported **tar** archive and load them. Red Hat makes the base image available in this form from <https://access.redhat.com/search/browse/docker-images#?>. With the **tar** archive in the current directory, the following command will load it into the local index:

```
[root@demo ~]# docker load -i rhel-server-docker-7.0-21.4-x86_64.tar.gz
```

Once the images are loaded into the local index, list the ones that are available with the command **docker images**:

```
[root@demo ~]# docker images
REPOSITORY                                TAG    IMAGE ID    CREATED    VIRTUAL SIZE
registry.access.redhat.com/redhat/rhel7    0      e1f5733f050b  3 days ago  140.2 MB
registry.access.redhat.com/redhat/rhel7    0-21   e1f5733f050b  3 days ago  140.2 MB
```

Running Docker containers

The **docker run** command is used to run a container. In its simplest form, it takes two arguments: the name of the image to use, and the name of a command to run in the container. Here is an example that will run a container in interactive mode, using the image **rhel17**, and run the command **cat /etc/hosts**. When the command exits, the container will be stopped. The **--rm** command will delete the container when it exits.

```
[root@demo ~]# docker run -i -t --rm rhel7 cat /etc/hosts
172.17.0.12 64fde39b4cea
127.0.0.1 localhost
::1 localhost ip6-localhost ip6-loopback
fe00::0 ip6-localnet
ff00::0 ip6-mcastprefix
ff02::1 ip6-allnodes
ff02::2 ip6-allrouters
```

A more useful example will start a new container, running the **bash** command:

```
[root@demo ~]# docker run -i -t rhel7 bash
bash-4.2# echo hello > /tmp/testfile
bash-4.2# exit
exit
[root@demo ~]#
```

This time, when the container exited, it did not remove its file system because the **--rm** option was not included. This can be seen with the **docker ps -a** command. The **-a** option shows all containers on the system, running or stopped. (Without **-a**, only running containers are displayed.)

```
[root@demo ~]# docker ps -a
CONTAINER ID  IMAGE          COMMAND  CREATED          STATUS          PORTS  NAMES
37917647b633  redhat/rhel7:0  bash    7 minutes ago   Exited (0) 2 minutes ago
silly_fermi3
```

Since the container was not removed when it exited, use **docker start** to restart it and connect to it interactively again (with the **-ai** options). The name of the container was assigned randomly, since one was not specified when it was first created. It can be selected by name or ID.

```
[root@demo ~]# docker start -ai silly_fermi3
bash-4.2# cat /tmp/testfile
hello
bash-4.2# exit
```




Important

The RHEL 7 base container image is an *extremely* minimal environment. Docker ensures it has network connectivity (despite the image not having networking utilities installed) and all it is designed to be able to do is use **yum** to install additional packages. (The container does not need a RHEL subscription in its own right, as it uses its host's entitlement.)

This is intentional. The idea behind the base container image is that it is the absolute minimum needed to allow the administrator to bootstrap it to a useful set of software packages for their purposes.

Creating Docker images

In the next example, a Docker image will be created from the base image that includes the *httpd* package. The new image will then be used to run a container.

Run a container from the base RHEL 7 image once, to manually install the *httpd* package.

```
[root@demo ~]# docker run -i rhel7 bash -c "yum install -y httpd; yum clean all"
```

After the packages have installed and the container has exited, identify the last container which was run with **docker ps -l**. With the name or ID of the container, the **docker commit** command can be used to create the new image from the container's file system. The following command will also include a comment (with **-m**).

```
[root@demo ~]# docker ps -l
CONTAINER ID   IMAGE          COMMAND                  CREATED        STATUS
PORTS         NAMES
ecaf8a3102a3  redhat/rhel7:0 bash -c 'yum clean a    4 minutes ago Exited (0) 3 minutes
ago          dreamy_goodall1
[root@demo ~]# docker commit -m "RHEL7 + httpd" dreamy_goodall1 rhel_httpd
74acf538ab6f5308a77bc308abae753e4ee683583eb4821455bccd4a141a7591
[root@demo ~]# docker images
REPOSITORY          TAG          IMAGE ID          CREATED          VIRTUAL SIZE
rhel_httpd          latest      74acf538ab6f    52 seconds ago  168.2 MB
registry.access.redhat.com/redhat/rhel7 0           e1f5733f050b    3 days ago     140.2 MB
registry.access.redhat.com/redhat/rhel7 0-21       e1f5733f050b    3 days ago     140.2 MB
```

Now a container can be run which uses this new image as its baseline. The following command will introduce two new options. The **-p 8080:80** option will map port 8080/TCP on the host to port 80/TCP on the container. The **-d** will cause the container to start in the background. The command **/usr/sbin/httpd -DFOREGROUND** will start Apache HTTPD in a very basic mode in the container.

```
[root@demo ~]# docker run -p 8080:80 -d rhel_httpd /usr/sbin/httpd -DFOREGROUND
de50d07b5a7e9e430558ec205fa49c0010adc5ccc19fb0d384081319925c10ab
[root@demo ~]# curl http://localhost:8080
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.1//EN" "http://www.w3.org/TR/xhtml11/DTD/
xhtml11.dtd">

<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en">
  <head>
    <title>Test Page for the Apache HTTP Server on Red Hat Enterprise Linux</title>

  ...
```

The `docker history image` command can be used to review the commit log for an image and to see what image layers it is based upon.



Note

An even better way to construct the image is to use `docker build` and a *Dockerfile* to script its construction. Like Kickstart, the advantage of using Dockerfiles is that the builds are reproducible at some point in the future. There is more information on how to do this in the references listed at the end of this section.

Starting a private Docker registry

To make it easier to distribute finished images to other hosts, system administrators can push their images to a Docker registry. Red Hat provides a package, *docker-registry*, which can be used to set up a private registry for an organization's internal use. This server does not need to be on a host running Docker itself.

The steps to set this up are simple:

1. Install *docker-registry* from RHEL 7 Extras.
2. Start and enable the `docker-registry` service.
3. Ensure that port 5000/TCP is reachable through the system's firewall.

```
[root@demo ~]# yum install docker-registry
...
[root@demo ~]# systemctl start docker-registry
[root@demo ~]# systemctl enable docker-registry
[root@demo ~]# firewall-cmd --add-port=5000/tcp
[root@demo ~]# firewall-cmd --add-port=5000/tcp --permanent
```

Pushing images to a private registry

To push an image to a private repository, it must first be tagged with the host address and port of the repository with `docker tag`. It can then be pushed with `docker push`.

```
[root@demo ~]# docker tag rhel_httpd registry.example.com:5000/user/httpd
[root@demo ~]# docker images
REPOSITORY                                TAG      IMAGE ID      CREATED        VIRTUAL SIZE
rhel_httpd                                latest   74acf538ab6f  36 minutes ago 168.2 MB
registry.example.com:5000/user/httpd      latest   74acf538ab6f  36 minutes ago 168.2 MB
registry.access.redhat.com/redhat/rhel7  0        e1f5733f050b  3 days ago    140.2 MB
registry.access.redhat.com/redhat/rhel7  0-21    e1f5733f050b  3 days ago    140.2 MB
[root@demo ~]# docker push registry.example.com:5000/user/httpd
The push refers to a repository [registry.example.com:5000/user/httpd] (len: 1)
Sending image list
Pushing repository registry.example.com:5000/user/httpd (1 tags)
Image e1f5733f050b already pushed, skipping
74acf538ab6f: Image successfully pushed
Pushing tag for rev [74acf538ab6f] on [http://registry.example.com:5000/v1/repositories/
user/httpd/tags/latest}
[root@demo ~]#
```

Exporting images to a file

The **docker save** command can be used to save an image file to an uncompressed **tar** archive. This file can then be copied to another Docker host and loaded into its index with **docker load**. This archive will include all the layers needed to reconstruct the image.

```
[root@demo ~]# docker save --output=rhel_httpd-latest.tar rhel_httpd
```

Cleaning up

Several commands are useful for cleaning up containers:

- **docker stop *container*** will stop a container gracefully, and **docker kill *container*** will send it a **kill** signal.
- **docker rm *container*** will delete a container's file system image permanently from the local index. Once this is done, the container cannot be restarted; it must be recreated.
- **docker rmi *image*** will remove an image from the host's index. It does not remove that image from a registry. It is important not to remove any images that a host may currently be using.
- **docker info** will provide some basic information about the Docker environment and current resource consumption.



References

docker(1) and **nsenter**(1) man pages

Get Started with Docker Containers in RHEL7
<https://access.redhat.com/site/articles/881893>

Red Hat Enterprise Linux 7 - Docker Overview
<https://access.redhat.com/site/videos/911053>

Docker
<http://www.docker.com>

More information is available in the "Linux Containers with Docker Format" section of the *Red Hat Enterprise Linux 7.0 Release Notes* and in the *Red Hat Enterprise Linux 7 Resource Management and Linux Containers Guide*, which can both be found at <https://access.redhat.com/docs>

Summary

Introduction to Linux Containers

In this section, students learned how to:

- Explain how Linux containers are implemented.
- Explain the role of Docker in container management.

Using Docker

In this section, students learned how to:

- Install and configure Docker on Red Hat Enterprise Linux 7.
- Create and run containers with Docker.
- Create new Docker images from a container.
- Install a private Docker registry and push images to it.



CHAPTER 15

COMPREHENSIVE REVIEW

Overview	
Goal	To practice and demonstrate knowledge and skills learned in Red Hat System Administration III.
Objectives	<ul style="list-style-type: none">• Review the course chapters to reinforce knowledge and skills.
Sections	<ul style="list-style-type: none">• Red Hat System Administration III Comprehensive Review
Lab	<ul style="list-style-type: none">• Comprehensive Review of System Administration III

Red Hat System Administration III

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 III

Before beginning the comprehensive review for this course, students should be comfortable with the topics covered in each chapter.

Do not hesitate to ask the instructor for extra guidance or clarification on these topics.

Chapter 1, Controlling Services and Daemons

To review how to manage services and the boot-up process using **systemctl**.

- Manage how **systemd** starts up system daemons and network services, using **systemctl**.
- Control and troubleshoot system boot using **systemd** targets.

Chapter 2, Managing IPv6 Networking

To configure and troubleshoot basic IPv6 networking on Red Hat Enterprise Linux systems.

- Review how to configure IPv4 networking in RHEL 7.
- Explain the basic concepts of IPv6 networking, and read and write condensed IPv6 addresses.
- Configure IPv6 networking using command-line tools and configuration files.

Chapter 3, Configuring Link Aggregation and Bridging

To configure and troubleshoot advanced network interface functionality, including teaming and local software bridges.

- Use network teaming to provide link redundancy or higher throughput.
- Manage a network team interface.
- Manage local software bridges and associated interfaces.

Chapter 4, Network Port Security

To permit and reject access to network services using advanced SELinux and firewalld filtering techniques.

- Review firewalld concepts and management commands covered in previous courses.
- Configure more complex firewall configurations using firewalld's support for "rich language rules."
- Describe and implement Network Address Translation (NAT).
- Ensure network ports have the correct SELinux type so that services are able to bind to them.

Chapter 5, Managing DNS for Servers

To set and verify correct DNS records for systems and configure secure caching DNS name service.

- Explain how DNS is used to resolve names and addresses and the purpose of key DNS resource records.
- Configure *unbound* to act as a secure local caching nameserver.
- Describe common DNS configuration problems and how to identify and resolve them.

Chapter 6, Configuring Email Transmission

To relay all email sent by the system through an SMTP gateway.

- Configure a Red Hat Enterprise Linux server to transmit all email through an unauthenticated SMTP gateway.

Chapter 7, Providing Remote Block Storage

To provide and use networked iSCSI block devices as remote disks.

- Explain at a high level how iSCSI is used to provide remote access to block devices.
- Provide remote access using a local disk as a LUN of an iSCSI storage target.
- Access remote storage using an iSCSI initiator and prepare it for use.

Chapter 8, Providing File-based Storage

To provide NFS exports and SMB file shares to specific systems and users.

- Export file systems to client systems using NFS, controlling access by IP address.
- Export file systems to clients using NFS, controlling access with Kerberos and using labeled NFS.
- Share file systems with clients using SMB, controlling access by username and password.
- Mount an SMB share with the **multiuser** mount option, using password-based authentication and **cifscreds** to control access

Chapter 9, Configuring MariaDB Databases

To provide a MariaDB SQL database for use by programs and database administrators.

- Install MariaDB.
- Configure and administer MariaDB.
- Configure user and access rights.
- Back up and restore MariaDB databases.

Chapter 10, Providing Apache HTTPD Web Service

To configure Apache HTTPD to provide TLS-enabled websites and virtual hosts.

- Identify the key configuration files, log files, and content directories used by Apache **httpd**.
- Configure Apache **httpd** to provide IP-based and name-based virtual hosts.
- Configure Apache **httpd** to provide TLS-encrypted virtual hosts.
- Configure Apache **httpd** to serve dynamic database-driven web content.

Chapter 11, Writing Bash Scripts

To write simple, well-structured shell scripts using Bash's shell expansion features and for-loop construct.

- To write simple shell scripts using Bash.

Chapter 12, Bash Conditionals and Control Structures

To use Bash conditionals and other control structures to write more sophisticated shell commands and scripts.

- Incorporate the use of positional parameters, exit status, test conditions, and conditional structures to implement flow control in Bash shell scripts.

Chapter 13, Configuring the Shell Environment

To customize Bash startup and use environment variables, Bash aliases, and Bash functions.

- Use bash startup scripts to define environment variables, aliases, and functions

Chapter 14, Linux Containers and Docker

To manage Linux containers with Docker in Red Hat Enterprise Linux 7.

- Explain what Linux containers are and Docker's role in managing them.
- Manage Linux containers using Docker.



References

Get information on more classes available from Red Hat at <http://www.redhat.com/training/>

Lab: Comprehensive Review of System Administration III

Task scenario

In this lab, you will configure a system using the skills taught in this course.

Resources:	
Files:	<ul style="list-style-type: none"> • <code>/home/student/myapp.php</code> • <code>/home/student/mydb.mysqldump</code>
Machines:	<code>serverX</code> and <code>desktopX</code>

Outcomes:

A system configured according to the listed specifications.

Before you begin...

- Reset your `desktopX` system.
- Log into and set up your `desktopX` system.

```
[student@desktopX ~]$ lab sa3-review setup
```

- Reset your `serverX` system.
- Log into and set up your `serverX` system.

```
[student@serverX ~]$ lab sa3-review setup
```

Specifications

You have been tasked with configuring a new system for your company: `serverX`. The system should be configured according to the following requirements:

- Your `serverX` system has two interfaces named `slave1` and `slave2`. These interfaces should be slaved in a new teaming device called `team1`.

The teaming device should use the `activebackup` runner, and be configured with the following IP addresses:

- `192.168.0.100/24`
- `fd00:ba5e:ba11:X::1/64`

When successfully configured, both `192.168.0.254` and `fd00:ba5e:ba11:X::fe` should be pingable.

- Your `serverX` system should provide caching DNS services to systems on the `172.25.X.0/24` subnet, using `172.25.254.254` as a forwarder.

DNSSEC validation for the `example.com` domain should be bypassed.

- Your **serverX** system should accept new email messages over SMTP from the **172.25.X.0/24** subnet. All messages not addressed to **@serverX.example.com** or **@localhost[.localdomain]** should be forwarded to the SMTP smarthost running on **desktopX.example.com**.
- Your **serverX** should export an iSCSI disk called **iqn.2014-06.com.example:serverX.zoidberg**. This iSCSI device should be backed by a **512 MiB** partition on the secondary hard drive. Access should only be allowed to clients with an IQN of **iqn.2014-06.com.example:desktopX**.
- Your **serverX** system should export the (new) **/exports/hypnotoad** directory using NFSv4. This export should require Kerberos authentication and full encryption. All clients on the **172.25.X.0/24** subnet should have read-write access to this export.

File system permissions on the underlying **/exports/hypnotoad** directory should be **1777**.

Both your **desktopX** and **serverX** systems have already been configured for LDAP and Kerberos for user information and authentication. **desktopX** has also been configured with a client keytab for Kerberos.

A keytab for your **serverX** system can be obtained from <http://classroom.example.com/pub/keytabs/serverX.keytab>.

- Your **serverX** system should export the (new) directory **/exports/bigbang** using **smb**. This directory should be owned by the **bigbang** group, and any new files created in this directory should automatically be owned by the **bigbang** group as well.

The name for the share should be **bigbang**, and the Samba-only user **penny** should be able to connect read-write using the password **I<3Sheldon**. **penny** should be a member of the **bigbang** group.

- Your **serverX** system should serve out two TLS-protected domains on port **444/TCP**: **wwwX.example.com** and **webappX.example.com**. These domains should serve out content from **/srv/wwwX/www/** and **/srv/webappX/www/**, respectively.

The TLS-enabled website on <https://wwwX.example.com:444> should respond with the following text:

This is **wwwX**

The website at <https://webappX.example.com:444> should respond by executing the PHP application found in **/home/student/myapp.php**.

You will need the following certificates and keys to complete this requirement:

- <http://classroom.example.com/pub/example-ca.crt>
- <http://classroom.example.com/pub/tls/certs/wwwX.crt>
- <http://classroom.example.com/pub/tls/private/wwwX.key>
- <http://classroom.example.com/pub/tls/certs/webappX.crt>
- <http://classroom.example.com/pub/tls/private/webappX.key>

The PHP application needs a database to work. A complete backup (**mysqldump**) of a working database can be found in **/home/student/mydb.mysqldump**.

- On your **desktopX** machine, create a new (executable) script called **/home/student/bin/myusers**.

When this script is called with the **userlist** argument, it should return a list of all users defined in **/etc/passwd** with a shell other than **/sbin/nologin**. Output should be a list of usernames, one per line, sorted alphabetically.

When this script is called with the **userinfo** argument, it should return the preferred shell of the second argument, or **invalid user** if that user does not exist. If the second argument is empty or undefined, the error message **“Please specify a username”** must be displayed, and an exit code of **132** must be returned.

If the script is called without any arguments, the following text must be displayed:

```
myusers userlist
myusers userinfo <USERNAME>
```

Examples:

```
[student@desktopX ~]$ myusers
myusers userlist
myusers userinfo <USERNAME>
[student@desktopX ~]$ myusers userlist
halt
root
shutdown
student
sync
[student@desktopX ~]$ myusers userinfo student
/bin/bash
```

Any error conditions not specified previously should return no output, and an exit code of **0**.

All your changes must survive a reboot. When you are done configuring your **serverX** system, reboot it, then run the following command on **desktopX** to verify your work.

```
[student@desktopX ~]$ lab sa3-review grade
```

Solution

In this lab, you will configure a system using the skills taught in this course.

Resources:	
Files:	<ul style="list-style-type: none"> • <code>/home/student/myapp.php</code> • <code>/home/student/mydb.mysqldump</code>
Machines:	serverX and desktopX

Outcomes:

A system configured according to the listed specifications.

Before you begin...

- Reset your **desktopX** system.
- Log into and set up your **desktopX** system.

```
[student@desktopX ~]$ lab sa3-review setup
```

- Reset your **serverX** system.
- Log into and set up your **serverX** system.

```
[student@serverX ~]$ lab sa3-review setup
```

Specifications

You have been tasked with configuring a new system for your company: **serverX**. The system should be configured according to the following requirements:

- Your **serverX** system has two interfaces named **slave1** and **slave2**. These interfaces should be slaved in a new teaming device called **team1**.

The teaming device should use the **activebackup** runner, and be configured with the following IP addresses:

- **192.168.0.100/24**
- **fd00:ba5e:ba11:X::1/64**

When successfully configured, both **192.168.0.254** and **fd00:ba5e:ba11:X::fe** should be pingable.

- Your **serverX** system should provide caching DNS services to systems on the **172.25.X.0/24** subnet, using **172.25.254.254** as a forwarder.

DNSSEC validation for the **example.com** domain should be bypassed.

- Your **serverX** system should accept new email messages over SMTP from the **172.25.X.0/24** subnet. All messages not addressed to **@serverX.example.com** or **@localhost[.localdomain]** should be forwarded to the SMTP smarthost running on **desktopX.example.com**.

- Your **serverX** should export an iSCSI disk called **iqn.2014-06.com.example:serverX.zoidberg**. This iSCSI device should be backed by a **512 MiB** partition on the secondary hard drive. Access should only be allowed to clients with an IQN of **iqn.2014-06.com.example:desktopX**.
- Your **serverX** system should export the (new) **/exports/hypnotoad** directory using NFSv4. This export should require Kerberos authentication and full encryption. All clients on the **172.25.X.0/24** subnet should have read-write access to this export.

File system permissions on the underlying **/exports/hypnotoad** directory should be **1777**.

Both your **desktopX** and **serverX** systems have already been configured for LDAP and Kerberos for user information and authentication. **desktopX** has also been configured with a client keytab for Kerberos.

A keytab for your **serverX** system can be obtained from <http://classroom.example.com/pub/keytabs/serverX.keytab>.

- Your **serverX** system should export the (new) directory **/exports/bigbang** using **smb**. This directory should be owned by the **bigbang** group, and any new files created in this directory should automatically be owned by the **bigbang** group as well.

The name for the share should be **bigbang**, and the Samba-only user **penny** should be able to connect read-write using the password **I<3Sheldon**. **penny** should be a member of the **bigbang** group.

- Your **serverX** system should serve out two TLS-protected domains on port **444/TCP**: **wwwX.example.com** and **webappX.example.com**. These domains should serve out content from **/srv/wwwX/www/** and **/srv/webappX/www/**, respectively.

The TLS-enabled website on <https://wwwX.example.com:444> should respond with the following text:

This is wwwX

The website at <https://webappX.example.com:444> should respond by executing the PHP application found in **/home/student/myapp.php**.

You will need the following certificates and keys to complete this requirement:

- <http://classroom.example.com/pub/example-ca.crt>
- <http://classroom.example.com/pub/tls/certs/wwwX.crt>
- <http://classroom.example.com/pub/tls/private/wwwX.key>
- <http://classroom.example.com/pub/tls/certs/webappX.crt>
- <http://classroom.example.com/pub/tls/private/webappX.key>

The PHP application needs a database to work. A complete backup (**mysqldump**) of a working database can be found in **/home/student/mydb.mysqldump**.

- On your **desktopX** machine, create a new (executable) script called **/home/student/bin/myusers**.

When this script is called with the **userlist** argument, it should return a list of all users defined in **/etc/passwd** with a shell other than **/sbin/nologin**. Output should be a list of usernames, one per line, sorted alphabetically.

When this script is called with the **userinfo** argument, it should return the preferred shell of the second argument, or **"invalid user"** if that user does not exist. If the second argument is empty or undefined, the error message **"Please specify a username"** must be displayed, and an exit code of **132** must be returned.

If the script is called without any arguments, the following text must be displayed:

```
myusers userlist
myusers userinfo <USERNAME>
```

Examples:

```
[student@desktopX ~]$ myusers
myusers userlist
myusers userinfo <USERNAME>
[student@desktopX ~]$ myusers userlist
halt
root
shutdown
student
sync
[student@desktopX ~]$ myusers userinfo student
/bin/bash
```

Any error conditions not specified previously should return no output, and an exit code of **0**.

All your changes must survive a reboot. When you are done configuring your **serverX** system, reboot it, then run the following command on **desktopX** to verify your work.

```
[student@desktopX ~]$ lab sa3-review grade
```

1. Your **serverX** system has two interfaces named **slave1** and **slave2**. These interfaces should be slaved in a new teaming device called **team1**.

The teaming device should use the **activebackup** runner, and be configured with the following IP addresses:

- **192.168.0.100/24**
- **fd00:ba5e:ba11:X::1/64**

When successfully configured, both **192.168.0.254** and **fd00:ba5e:ba11:X::fe** should be pingable.

1. Create a new teaming device called **team1**, using the **activebackup** runner.

```
[student@serverX ~]$ sudo nmcli con add type team con-name team1 ifname team1
config '{"runner": {"name": "activebackup"}}'
```

2. Add the IP addresses to the **team1** config.

```
[student@serverX ~]$ sudo nmcli con mod team1 ipv4.addresses '192.168.0.100/24'
[student@serverX ~]$ sudo nmcli con mod team1 ipv4.method static
[student@serverX ~]$ sudo nmcli con mod team1 ipv6.addresses
'fd00:ba5e:ba11:X::1/64'
[student@serverX ~]$ sudo nmcli con mod team1 ipv6.method static
```

3. Add both the **slave1** and **slave2** interfaces as slaves to the **team1** interface.

```
[student@serverX ~]$ sudo nmcli con add type team-slave con-name team1-slave1
ifname slave1 master team1
[student@serverX ~]$ sudo nmcli con add type team-slave con-name team1-slave2
ifname slave2 master team1
```

4. Test by pinging both **192.168.0.254** and **fd00:ba5e:ba11:X::fe**.

```
[student@serverX ~]$ ping -c1 192.168.0.254
[student@serverX ~]$ ping6 -c1 fd00:ba5e:ba11:X::fe
```

2. Your **serverX** system should provide caching DNS services to systems on the **172.25.X.0/24** subnet, using **172.25.254.254** as a forwarder.

DNSSEC validation for the **example.com** domain should be bypassed.

1. Install *unbound* on your **serverX** system.

```
[student@serverX ~]$ sudo yum install unbound
```

2. Configure **unbound** to listen on all interfaces. Add the following lines to **/etc/unbound/unbound.conf** inside the **server:** block:

```
interface: 0.0.0.0
interface: ::0
```

3. Allow queries from the **172.25.X.0/24** subnet. Add the following line to **/etc/unbound/unbound.conf** inside the **server:** block:

```
access-control: 172.25.X.0/24 allow
```

4. Disable DNSSEC validation for the **example.com** domain. Add the following line to **/etc/unbound/unbound.conf** inside the **server:** block:

```
domain-insecure: example.com
```


5. Configure **unbound** to forward all requests to **172.25.254.254**. Add the following lines to **/etc/unbound/unbound.conf**:

```
forward-zone:
  name: .
  forward-addr: 172.25.254.254
```

6. Start and enable the **unbound.service** service.

```
[student@serverX ~]$ sudo systemctl enable unbound.service
[student@serverX ~]$ sudo systemctl start unbound.service
```

7. Configure the default firewall zone to allow DNS traffic.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --add-service=dns
[student@serverX ~]$ sudo firewall-cmd --reload
```

3. Your **serverX** system should accept new email messages over SMTP from the **172.25.X.0/24** subnet. All messages not addressed to **@serverX.example.com** or **@localhost.localdomain** should be forwarded to the SMTP smarthost running on **desktopX.example.com**.

1. Verify that the **postfix** package is installed on **serverX**.

```
[student@serverX ~]$ sudo yum install postfix
```

2. Configure **postfix** to listen on all network interfaces, relaying mail from the **172.25.X.0/24** (and localhost) subnet.

```
[student@serverX ~]$ sudo postconf -e 'inet_interfaces = all'
[student@serverX ~]$ sudo postconf -e 'mynetworks = 172.25.X.0/24, 127.0.0.0/8'
```

3. Configure **postfix** to only perform local delivery for the **serverX.example.com**, **localhost**, and **localhost.localdomain** domains.

```
[student@serverX ~]$ sudo postconf -e 'mydestination = serverX.example.com,
localhost.localdomain, localhost'
```

4. Configure **postfix** to send all mail not included in **mydestination** to the server at **desktopX.example.com**.

```
[student@serverX ~]$ sudo postconf -e 'relayhost = [desktopX.example.com]'
```

5. Open all relevant ports on the firewall.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --add-service=smtp
[student@serverX ~]$ sudo firewall-cmd --reload
```

- (Re)start and enable the **postfix.service** service.

```
[student@serverX ~]$ sudo systemctl restart postfix.service
[student@serverX ~]$ sudo systemctl enable postfix.service
```

- Create a new **512 MiB** partition on your secondary hard drive on **serverX**, then export that over iSCSI using the name **iqn.2014-06.com.example:serverX.zoidberg**, but only to clients with the **iqn.2014-06.com.example:desktopX** initiator name.

- Install the necessary packages.

```
[student@serverX ~]$ sudo yum install targetd targetcli
```

- Start and enable the **targetd.service** service.

```
[student@serverX ~]$ sudo systemctl start targetd.service
[student@serverX ~]$ sudo systemctl enable targetd.service
```

- Create a new **512 MiB** partition on your secondary hard drive.

```
[student@serverX ~]$ sudo fdisk /dev/vdb
Welcome to fdisk (util-linux 2.23.2).

Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Device does not contain a recognized partition table
Building a new DOS disklabel with disk identifier 0x0aee5076.

Command (m for help): n
Partition type:
   p   primary (0 primary, 0 extended, 4 free)
   e   extended
Select (default p): p
Partition number (1-4, default 1): 1
First sector (2048-20971519, default 2048): Enter
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519): +512M
Partition 1 of type Linux and of size 512 MiB is set

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

- Configure **targetd** for the target, including the ACL. Make sure the configuration persists. Remember to use the **Tab** key liberally.

```
[student@serverX ~]$ sudo targetcli
targetcli shell version 2.1.fb34
Copyright 2011-2013 by Datera, Inc. and others.
For help on commands, type 'help'

/> backstores/block create serverX.zoidberg /dev/vdb1
/> /iscsi create iqn.2014-06.com.example:serverX.zoidberg
/> /iscsi/iqn.2014-06.com.example:serverX.zoidberg/tpg1/luns create /backstores/
block/serverX.zoidberg 1
/> /iscsi/iqn.2014-06.com.example:serverX.zoidberg/tpg1/portals create
172.25.X.11 3260
/> /iscsi/iqn.2014-06.com.example:serverX.zoidberg/tpg1/acls create
iqn.2014-06.com.example:desktopX
/> saveconfig
/> exit
```

5. Open port **3260/TCP** on the firewall.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --add-port=3260/tcp
[student@serverX ~]$ sudo firewall-cmd --reload
```

5. Configure a Kerberized NFSv4 export for **/exports/hypnotoad**. A keytab can be obtained from <http://classroom.example.com/pub/keytabs/serverX.keytab>. The export should be available read-write to clients using both Kerberos authentication and encryption. The exported directory should have **1777** permissions.

1. Create the exported directory.

```
[student@serverX ~]$ sudo mkdir -m1777 -p /exports/hypnotoad
```

2. Create a new file **/etc/exports** with the following contents:

```
/exports/hypnotoad 172.25.X.0/24(sec=krb5p,rw)
```

3. Download the keytab to **/etc/krb5.keytab**:

```
[student@serverX ~]$ sudo wget -O /etc/krb5.keytab http://classroom.example.com/
pub/keytabs/serverX.keytab
```

4. Start and enable the **nfs-secure-server.service** service.

```
[student@serverX ~]$ sudo systemctl start nfs-secure-server.service
[student@serverX ~]$ sudo systemctl enable nfs-secure-server.service
```

5. Open the ports needed for NFSv4 on the firewall.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --add-service=nfs
[student@serverX ~]$ sudo firewall-cmd --reload
```

6. Your **serverX** system should export the (new) directory **/exports/bigbang** using **smb**. This directory should be owned by the **bigbang** group, and any new files created in this directory should automatically be owned by the **bigbang** group as well. Users outside of the **bigbang** group should not have access to this share.

The name for the share should be **bigbang**, and the Samba-only user **penny** should be able to connect read-write using the password **I<3Sheldon**. **penny** should be a member of the **bigbang** group.

1. Create the user and group.

```
[student@serverX ~]$ sudo groupadd bigbang
[student@serverX ~]$ sudo useradd -s /sbin/nologin -G bigbang penny
```

2. Create the **/exports/bigbang** directory with permissions **2770**, owned by the **bigbang** group.

```
[student@serverX ~]$ sudo mkdir -m 2770 -p /exports/bigbang
[student@serverX ~]$ sudo chgrp bigbang /exports/bigbang
```

3. Mark the **/exports/bigbang** directory, and all of its children, as being of the SELinux **samba_share_t** context.

```
[student@serverX ~]$ sudo semanage fcontext -a -t samba_share_t '/exports/
bigbang(/.*)?'
[student@serverX ~]$ sudo restorecon -RFv /exports/bigbang
```

4. Install both the **samba** and **samba-client** packages.

```
[student@serverX ~]$ sudo yum install samba samba-client
```

5. Add an export block for the **/exports/bigbang** directory to the **/etc/samba/smb.conf** file. Use the share name **bigbang**.

```
[bigbang]
path = /exports/bigbang
valid users = @bigbang
write list = @bigbang
```

6. Set the Samba password for **penny** to **I<3Sheldon**.

```
[student@serverX ~]$ sudo smbpasswd -a penny
New SMB password: I<3Sheldon
Retype new SMB password: I<3Sheldon
Added user penny.
```

7. Start and enable the **smb.service** service.

```
[student@serverX ~]$ sudo systemctl start smb.service
[student@serverX ~]$ sudo systemctl enable smb.service
```

- Open up all the relevant ports on the firewall.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --add-service=samba
[student@serverX ~]$ sudo firewall-cmd --reload
```

- Your **serverX** system should serve out two TLS-protected domains on port **444/TCP**: **wwwX.example.com** and **webappX.example.com**. These domains should serve out content from **/srv/wwwX/www/** and **/srv/webappX/www/**, respectively.

The TLS-enabled website on **https://wwwX.example.com:444** should respond with the following text:

```
This is wwwX
```

The website at **https://webappX.example.com:444** should respond by executing the PHP application found in **/home/student/myapp.php**.

You will need the following certificates and keys to complete this requirement:

- http://classroom.example.com/pub/example-ca.crt**
- http://classroom.example.com/pub/tls/certs/wwwX.crt**
- http://classroom.example.com/pub/tls/private/wwwX.key**
- http://classroom.example.com/pub/tls/certs/webappX.crt**
- http://classroom.example.com/pub/tls/private/webappX.key**

The PHP application needs a database to work. A complete backup (**mysqldump**) of a working database can be found in **/home/student/mydb.mysqldump**.

- Install the needed packages:

```
[student@serverX ~]$ sudo yum install httpd mod_ssl mariadb-server php php-mysql
```

- Start and enable the **mariadb** service, then restore the database dump.

```
[student@serverX ~]$ sudo systemctl start mariadb
[student@serverX ~]$ sudo systemctl enable mariadb
[student@serverX ~]$ sudo mysql < /home/student/mydb.mysqldump
```

- Create the content directories and their content.

```
[student@serverX ~]$ sudo mkdir -p /srv/{www,webapp}X/www
[student@serverX ~]$ sudo cp /home/student/myapp.php /srv/webappX/www/index.php
[student@serverX ~]$ echo "This is wwwX" | sudo tee /srv/wwwX/www/index.html
[student@serverX ~]$ sudo restorecon -RFV /srv/
```

- Download all needed certificates.

```
[student@serverX ~]$ cd /etc/pki/tls/certs
[student@serverX certs]$ sudo wget http://classroom.example.com/pub/example-
ca.crt
[student@serverX certs]$ sudo wget http://classroom.example.com/pub/tls/certs/
wwwX.crt
[student@serverX certs]$ sudo wget http://classroom.example.com/pub/tls/certs/
webappX.crt
[student@serverX certs]$ cd /etc/pki/tls/private
[student@serverX private]$ sudo wget http://classroom.example.com/pub/tls/
private/wwwX.key
[student@serverX private]$ sudo wget http://classroom.example.com/pub/tls/
private/webappX.key
[student@serverX private]$ sudo chmod 0400 w*key
```

5. Configure SELinux to allow **httpd** to listen on port **444/TCP**.

```
[student@serverX ~]$ sudo semanage port -a -t http_port_t -p tcp 444
```

6. Configure **httpd** to listen on port **444/TCP** using https by adding the following line to **/etc/httpd/conf/httpd.conf**:

```
Listen 444 https
```

7. Configure two virtual hosts by adding the following two files to **/etc/httpd/conf.d/**:

/etc/httpd/conf.d/wwwX.conf

```
<VirtualHost *:444>
  ServerName wwwX.example.com
  SSLEngine On
  SSLProtocol all -SSLv2 -SSLv3
  SSLCipherSuite HIGH:MEDIUM:!aNULL:!MD5
  SSLHonorCipherOrder on
  SSLCertificateFile /etc/pki/tls/certs/wwwX.crt
  SSLCertificateKeyFile /etc/pki/tls/private/wwwX.key
  SSLCertificateChainFile /etc/pki/tls/certs/example-ca.crt
  DocumentRoot /srv/wwwX/www
</VirtualHost>

<Directory /srv/wwwX/www>
  Require all granted
</Directory>
```

/etc/httpd/conf.d/webappX.conf

```
<VirtualHost *:444>
  ServerName webappX.example.com
  SSLEngine On
  SSLProtocol all -SSLv2 -SSLv3
  SSLCipherSuite HIGH:MEDIUM:!aNull:!MD5
  SSLHonorCipherOrder on
  SSLCertificateFile /etc/pki/tls/certs/webappX.crt
  SSLCertificateKeyFile /etc/pki/tls/private/webappX.key
  SSLCertificateChainFile /etc/pki/tls/certs/example-ca.crt
  DocumentRoot /srv/webappX/www
</VirtualHost>

<Directory /srv/webappX/www>
  Require all granted
</Directory>
```

8. Start and enable the **httpd.service** service.

```
[student@serverX ~]$ sudo systemctl start httpd.service
[student@serverX ~]$ sudo systemctl enable httpd.service
```

9. Open all needed ports on the firewall.

```
[student@serverX ~]$ sudo firewall-cmd --permanent --add-port=444/tcp
[student@serverX ~]$ sudo firewall-cmd --reload
```

8. On your **desktopX** machine, create a new (executable) script called **/home/student/bin/myusers**.

When this script is called with the **userlist** argument, it should return a list of all users defined in **/etc/passwd**, with a shell other than **/sbin/nologin**. Output should be a list of usernames, one per line, sorted alphabetically.

When this script is called with the **userinfo** argument, it should return the preferred shell of the second argument, or **"invalid user"** if that user does not exist. If the second argument is empty or undefined, the error message **"Please specify a username"** must be displayed, and an exit code of **132** must be returned.

If the script is called without any arguments, the following text must be displayed:

```
myusers userlist
myusers userinfo <USERNAME>
```

Examples:

```
[student@serverX ~]$ myusers
myusers userlist
myusers userinfo <USERNAME>
[student@serverX ~]$ myusers userlist
halt
root
shutdown
student
sync
[student@serverX ~]$ myusers userinfo student
/bin/bash
```

Any error conditions not previously specified should return no output, and an exit code of **0**.

1. Create the **/home/student/bin/** directory:

```
[student@desktopX ~]$ mkdir /home/student/bin
```

2. Create the new file **/home/student/bin/myusers** with the following content:

```
#!/bin/bash

if [ $# -eq 0 ]; then
    echo "$(basename $0) userlist"
    echo "$(basename $0) userinfo <USERNAME>"
fi

case $1 in
    userlist)
        grep -v ':/sbin/nologin$' /etc/passwd | cut -d: -f1 | sort
        ;;
    userinfo)
        if [ "$2" == "" ]; then
            echo "Please specify a username"
            exit 132
        fi
        if ! getent passwd $2 &> /dev/null; then
            echo "invalid user"
            exit
        fi
        getent passwd $2 | cut -d: -f7
        ;;
    *)
        exit
        ;;
esac
```

3. Make the new script executable:

```
[student@desktopX ~]$ chmod +x /home/student/bin/myusers
```

9. Reboot your **serverX** machine, then verify your work by running the **lab sa3-review grade** command on your **desktopX** system.

- ```
[student@desktopX ~]$ lab sa3-review grade
```

## Summary

Red Hat System Administration III Comprehensive Review

- Review chapters to validate knowledge level.
- Review practice exercises to validate skill level.

