CompTIA.XK0-005.vMar-2024.by.Toan.94q

Number: XK0-005
Passing Score: 800
Time Limit: 120
File Version: 5.5

Exam Code: XK0-005

Exam Name: CompTIA Linux+

Exam A

QUESTION 1
A Linux system fails to start and delivers the following error message:

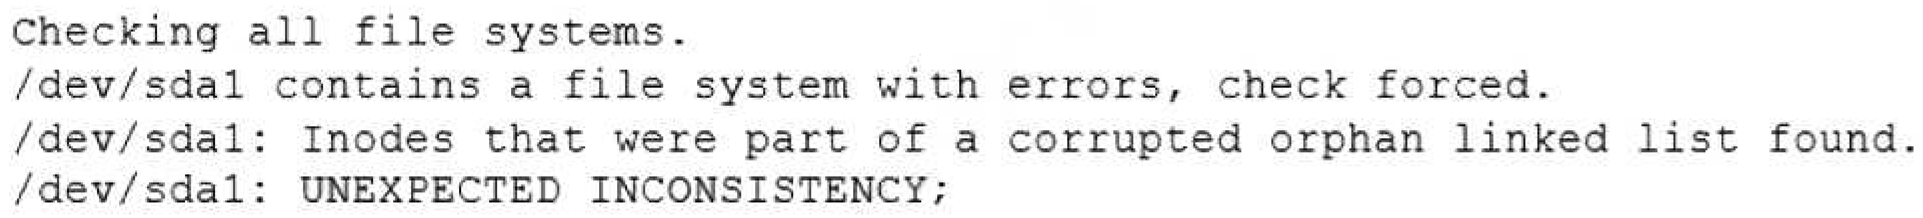

```
Checking all file systems.
/dev/sda1 contains a file system with errors, check forced.
/dev/sda1: Inodes that were part of a corrupted orphan linked list found.
/dev/sda1: UNEXPECTED INCONSISTENCY;
```

Which of the following commands can be used to address this issue?

A. fsck.ext4 /dev/sda1
B. partprobe /dev/sda1
C. fdisk /dev/sda1
D. mkfs.ext4 /dev/sda1

Correct Answer: A
Section:
Explanation:
The command fsck.ext4 /dev/sda1 can be used to address the issue. The issue is caused by a corrupted filesystem on the /dev/sda1 partition. The error message shows that the filesystem type is ext4 and the superblock is invalid. The command fsck.ext4 is a tool for checking and repairing ext4 filesystems. The command will scan the partition for errors and attempt to fix them. This command can resolve the issue and allow the system to start. The other options are incorrect because they either do not fix the filesystem (partprobe or fdisk) or destroy the data on the partition (mkfs.ext4). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 325.

QUESTION 2
Based on an organization's new cybersecurity policies, an administrator has been instructed to ensure that, by default, all new users and groups that are created fall within the specified values below.

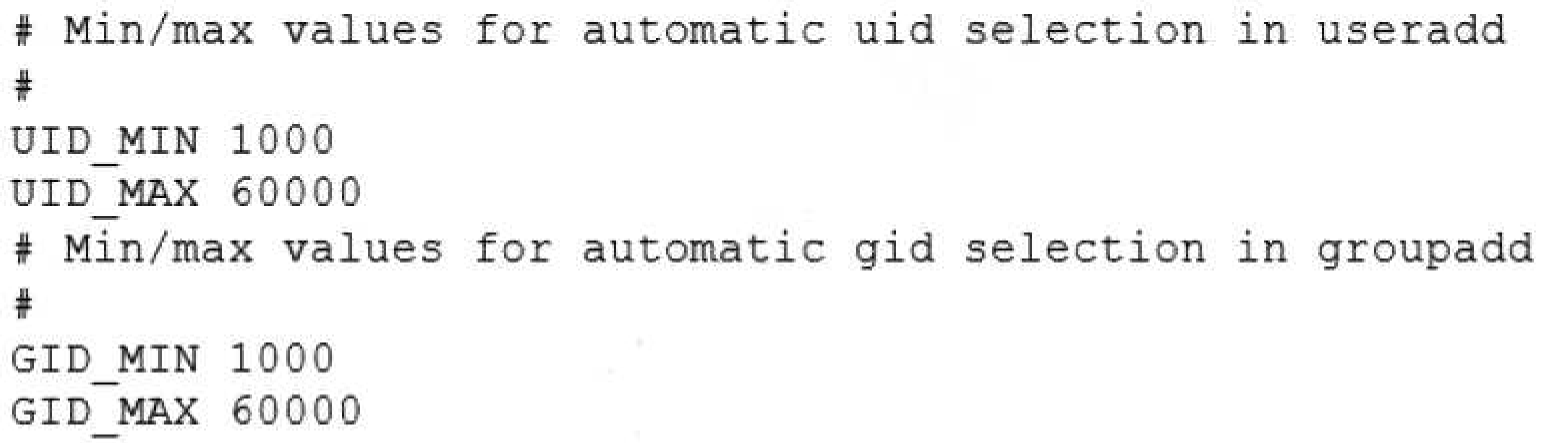

```
# Min/max values for automatic uid selection in useradd
#
UID_MIN 1000
UID_MAX 60000
# Min/max values for automatic gid selection in groupadd
#
GID_MIN 1000
GID_MAX 60000
```

To which of the following configuration files will the required changes need to be made?

A. /etc/login.defs
B. /etc/security/limits.conf

C. /etc/default/useradd

D. /etc/profile

Correct Answer: A
Section:
Explanation:
The required changes need to be made to the /etc/login.defs configuration file. The /etc/login.defs file defines the default values for user and group IDs, passwords, shells, and other parameters for user and group creation. The file contains the directives UID_MIN, UID_MAX, GID_MIN, and GID_MAX, which set the minimum and maximum values for automatic user and group ID selection. The administrator can edit this file and change the values to match the organization's new cybersecurity policies. This is the correct file to modify to accomplish the task. The other options are incorrect because they either do not affect the user and group IDs (/etc/security/limits.conf or /etc/profile) or do not set the default values (/etc/default/useradd). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15:
Managing Users and Groups, page 463.

QUESTION 3

A Linux administrator is trying to remove the ACL from the file /home/user/dat a. txt but receives the following error message:

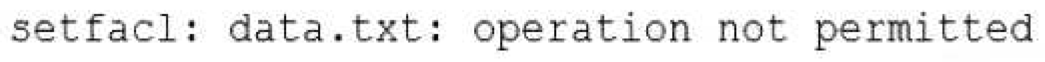

```
setfacl: data.txt: operation not permitted
```

Given the following analysis:

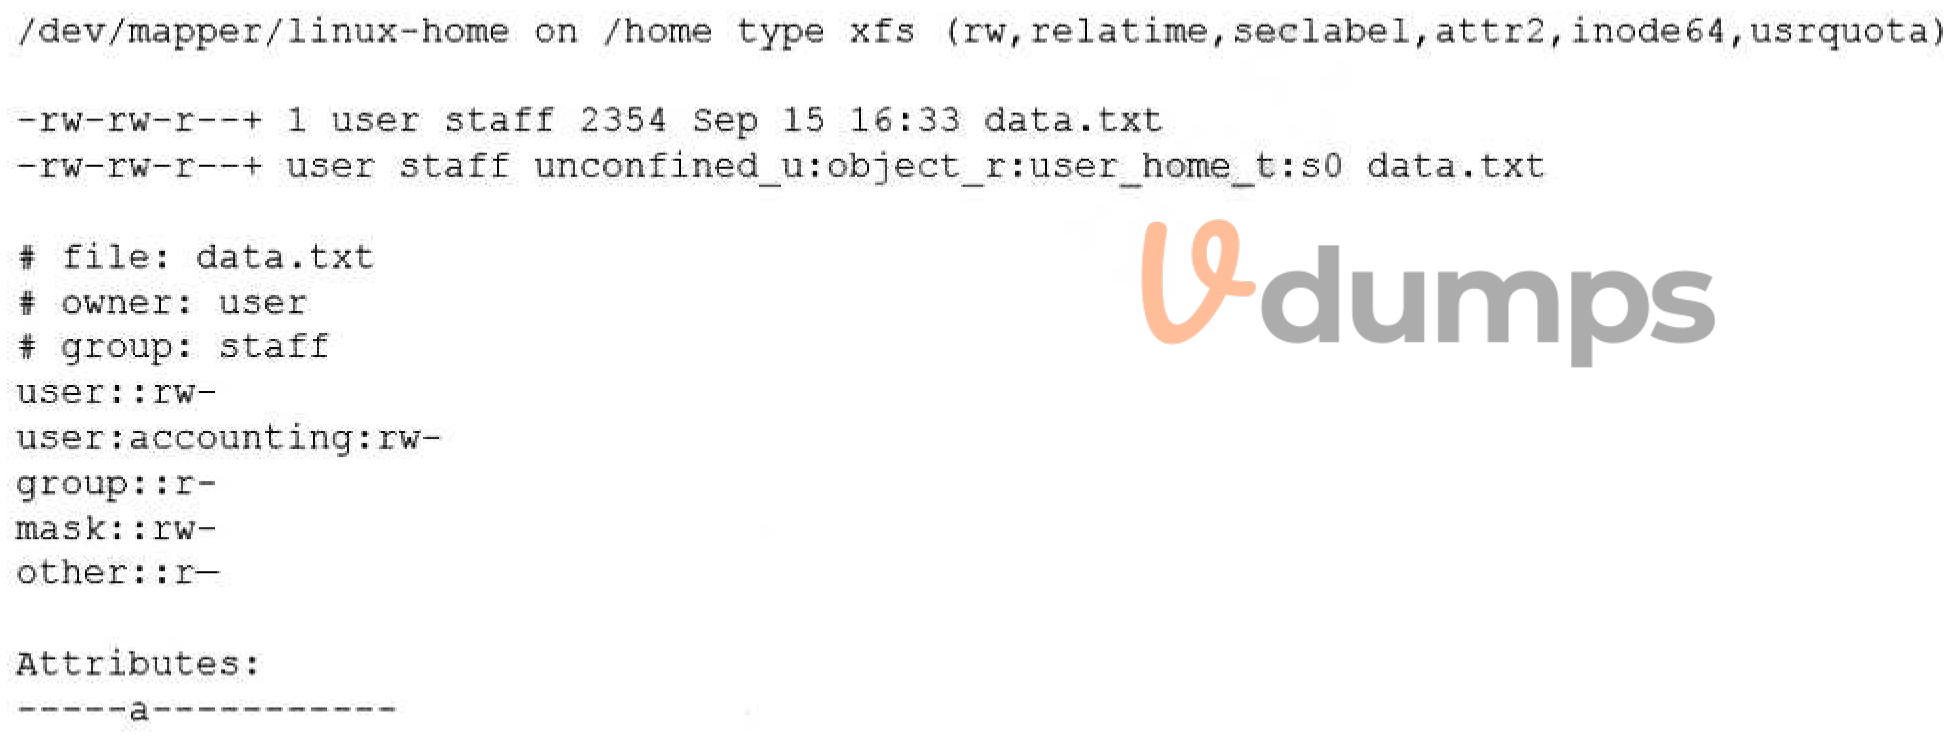

```
/dev/mapper/linux-home on /home type xfs (rw,relatime,seclabel,attr2,inode64,usrquota)

-rw-rw-r--+ 1 user staff 2354 Sep 15 16:33 data.txt
-rw-rw-r--+ user staff unconfined_u:object_r:user_home_t:s0 data.txt

# file: data.txt
# owner: user
# group: staff
user::rw-
user:accounting:rw-
group::r-
mask::rw-
other::r–

Attributes:
-----a-----------
```

Which of the following is causing the error message?

A. The administrator is not using a highly privileged account.

B. The filesystem is mounted with the wrong options.

C. SELinux file context is denying the ACL changes.

D. File attributes are preventing file modification.

Correct Answer: D
Section:
Explanation:
File attributes are preventing file modification, which is causing the error message. The output of lsattr /home/user/data.txt shows that the file has the immutable attribute (i) set, which means that the file cannot be changed, deleted, or renamed. The command setfacl -b /home/user/data.txt tries to remove the ACL from the file, but fails because of the immutable attribute. The administrator needs to remove the immutable attribute

first by using the command chattr -i /home/user/data.txt and then try to remove the ACL again. The other options are incorrect because they are not supported by the outputs. The administrator is using a highly privileged account, as shown by the # prompt. The filesystem is mounted with the correct options, as shown by the output of mount | grep /home. SELinux file context is not denying the ACL changes, as shown by the output of ls -Z /home/user/data.txt. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, pages 357-358.

QUESTION 4
A Linux administrator needs to create a new cloud.cpio archive containing all the files from the current directory. Which of the following commands can help to accomplish this task?

A. ls | cpio -iv > cloud.epio
B. ls | cpio -iv < cloud.epio
C. ls | cpio -ov > cloud.cpio
D. ls cpio -ov < cloud.cpio

Correct Answer: C
Section:
Explanation:
The command ls | cpio -ov > cloud.cpio can help to create a new cloud.cpio archive containing all the files from the current directory. The ls command lists the files in the current directory and outputs them to the standard output. The | operator pipes the output to the next command.
The cpio command is a tool for creating and extracting compressed archives. The -o option creates a new archive and the -v option shows the verbose output. The > operator redirects the output to the cloud.cpio file. This command will create a new cloud.cpio archive with all the files from the current directory. The other options are incorrect because they either use the wrong options (-i instead of -o), the wrong arguments (cloud.epio instead of cloud.cpio), or the wrong syntax (< instead of > or missing |). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, page 351.

QUESTION 5
A systems administrator made some changes in the ~/.bashrc file and added an alias command.
When the administrator tried to use the alias command, it did not work. Which of the following should be executed FIRST?

A. source ~/.bashrc
B. read ~/.bashrc
C. touch ~/.bashrc
D. echo ~/.bashrc

Correct Answer: A
Section:
Explanation:
The command source ~/.bashrc should be executed first to use the alias command.
The source command reads and executes commands from a file in the current shell environment.
The ~/.bashrc file is a configuration file that contains commands and aliases that are executed when a new bash shell is started. The administrator made some changes in the ~/.bashrc file and added an alias command, but the changes are not effective until the file is sourced or a new shell is started.
The command source ~/.bashrc will reload the file and make the alias command available. The other options are incorrect because they either do not execute the commands in the file (read, touch, or echo) or do not affect the current shell environment (read or echo). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 9: Working with the Linux Shell, page 295.

QUESTION 6
A junior systems administrator has just generated public and private authentication keys for passwordless login. Which of the following files will be moved to the remote servers?

A. id_dsa.pem
B. id_rsa
C. id_ecdsa
D. id_rsa.pub

Correct Answer: D
Section:
Explanation:
The file id_rsa.pub will be moved to the remote servers for passwordless login. The id_rsa.pub file is the public authentication key that is generated by the ssh-keygen command. The public key can be copied to the remote servers by using the ssh-copy-id command or manually. The remote servers will use the public key to authenticate the user who has the corresponding private key (id_rsa). This will allow the user to log in without entering a password. The other options are incorrect because they are either private keys (id_rsa, id_dsa.pem, or id_ecdsa) or non-existent files (id_dsa.pem or id_ecdsa). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 410.

QUESTION 7
An administrator accidentally deleted the /boot/vmlinuz file and must resolve the issue before the server is rebooted. Which of the following commands should the administrator use to identify the correct version of this file?

A. rpm -qa | grep kernel; uname -a
B. yum -y update; shutdown -r now
C. cat /etc/centos-release; rpm -Uvh --nodeps
D. telinit 1; restorecon -Rv /boot

Correct Answer: A
Section:
Explanation:
The command rpm -qa | grep kernel lists all the installed kernel packages, and the command uname -a displays the current kernel version. These commands can help the administrator identify the correct version of the /boot/vmlinuz file, which is the kernel image file. The other options are not relevant or helpful for this task. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing the Linux Boot Process, page 267.

QUESTION 8
A cloud engineer needs to change the secure remote login port from 22 to 49000. Which of the following files should the engineer modify to change the port number to the desired value?

A. /etc/host.conf
B. /etc/hostname
C. /etc/services
D. /etc/ssh/sshd_config

Correct Answer: D
Section:
Explanation:
The file /etc/ssh/sshd_config contains the configuration settings for the SSH daemon, which handles the secure remote login. To change the port number, the engineer should edit this file and modify the line that says Port 22 to Port 49000. The other files are not related to the SSH service. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 411.

QUESTION 9
A new file was added to a main Git repository. An administrator wants to synchronize a local copy with the contents of the main repository. Which of the following commands should the administrator use for this task?

A. git reflog
B. git pull
C. git status
D. git push

Correct Answer: B
Section:
Explanation:

The command iptables -t nat -A PREROUTING -p tcp --dport 80 -j DNAT --to-destination 192.0.2.25:3128 adds a rule to the nat table that redirects all incoming TCP packets with destination port 80 (HTTP) to the proxy server 192.0.2.25 on port 3128. This is the correct way to achieve the task. The other options are incorrect because they either delete a rule (-D), use the wrong protocol (top instead of tcp), or use the wrong port (81 instead of 80). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 381.

QUESTION 10
A Linux administrator needs to redirect all HTTP traffic temporarily to the new proxy server 192.0.2.25 on port 3128. Which of the following commands will accomplish this task?

A. iptables -t nat -D PREROUTING -p tcp --sport 80 -j DNAT - -to-destination 192.0.2.25:3128

B. iptables -t nat -A PREROUTING -p top --dport 81 -j DNAT --to-destination 192.0.2.25:3129

C. iptables -t nat -I PREROUTING -p top --sport 80 -j DNAT --to-destination 192.0.2.25:3129

D. iptables -t nat -A PREROUTING -p tcp --dport 80 -j DNAT --to-destination 192.0.2.25:3128

Correct Answer: D
Section:
Explanation:
The command iptables -t nat -A PREROUTING -p tcp --dport 80 -j DNAT --to-destination 192.0.2.25:3128 adds a rule to the nat table that redirects all incoming TCP packets with destination port 80 (HTTP) to the proxy server 192.0.2.25 on port 3128. This is the correct way to achieve the task. The other options are incorrect because they either delete a rule (-D), use the wrong protocol (top instead of tcp), or use the wrong port (81 instead of 80). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 381.

QUESTION 11
Developers have requested implementation of a persistent, static route on the application server.
Packets sent over the interface eth0 to 10.0.213.5/32 should be routed via 10.0.5.1. Which of the following commands should the administrator run to achieve this goal?

A. route -i etho -p add 10.0.213.5 10.0.5.1

B. route modify eth0 +ipv4.routes "10.0.213.5/32 10.0.5.1"

C. echo "10.0.213.5 10.0.5.1 eth0" > /proc/net/route

D. ip route add 10.0.213.5/32 via 10.0.5.1 dev eth0

Correct Answer: D
Section:
Explanation:
The command ip route add 10.0.213.5/32 via 10.0.5.1 dev eth0 adds a static route to the routing table that sends packets destined for 10.0.213.5/32 (a single host) through the gateway 10.0.5.1 on the interface eth0. This is the correct way to achieve the goal. The other options are incorrect because they either use the wrong syntax (route -i etho -p add), the wrong command (route modify), or the wrong file (/proc/net/route). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 379.

QUESTION 12
A user is asking the systems administrator for assistance with writing a script to verify whether a file exists. Given the following:

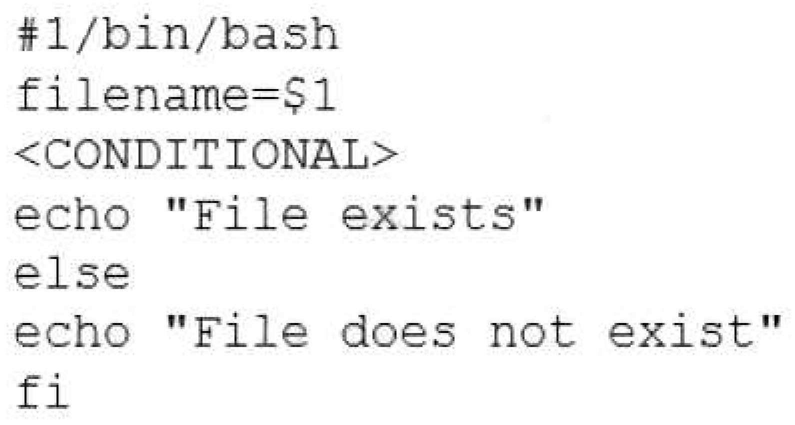

```
#1/bin/bash
filename=$1
<CONDITIONAL>
echo "File exists"
else
echo "File does not exist"
fi
```

Which of the following commands should replace the <CONDITIONAL> string?

A. if [-f "$filename"]; then
B. if [-d "$filename"]; then
C. if [-f "$filename"] then
D. if [-f "$filename"]; while

Correct Answer: A
Section:
Explanation:
The command if [-f "$filename"]; then checks if the variable $filename refers to a regular file that exists. The -f option is used to test for files. If the condition is true, the commands after then are executed. This is the correct way to replace the <CONDITIONAL> string. The other options are incorrect because they either use the wrong option (-d tests for directories), the wrong syntax (missing a semicolon after the condition), or the wrong keyword (while is used for loops, not conditions). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Writing and Executing Bash Shell Scripts, page 493.

QUESTION 13
A systems administrator is deploying three identical, cloud-based servers. The administrator is using the following code to complete the task:

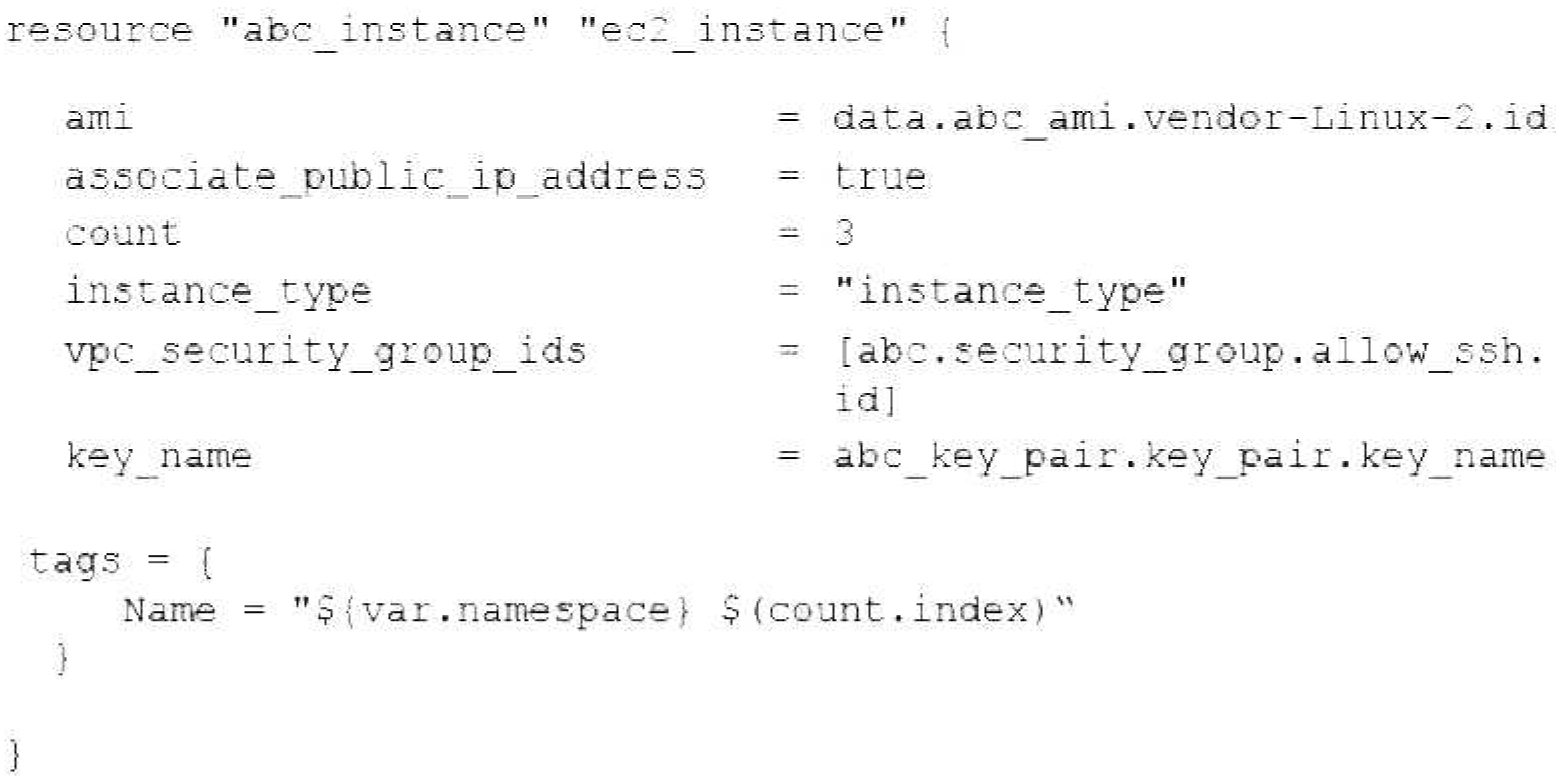

```
resource "abc_instance" "ec2_instance" {

  ami                          = data.abc_ami.vendor-Linux-2.id
  associate_public_ip_address  = true
  count                        = 3
  instance_type                = "instance_type"
  vpc_security_group_ids       = [abc.security_group.allow_ssh.
                                   id]
  key_name                     = abc_key_pair.key_pair.key_name

 tags = {
     Name = "${var.namespace} $(count.index)"
  }

}
```

Which of the following technologies is the administrator using?

A. Ansible
B. Puppet
C. Chef
D. Terraform

Correct Answer: D
Section:
Explanation:
The code snippet is written in Terraform language, which is a tool for building, changing, and versioning infrastructure as code. Terraform uses a declarative syntax to describe the desired state of the infrastructure and applies the changes accordingly. The code defines a resource of type aws_instance, which creates an AWS EC2 instance, and sets the attributes such as the AMI ID, instance type, security group IDs, and key name. The code also uses a count parameter to create three identical instances and assigns them different names using the count.index variable. This is the correct technology that the administrator is using. The other options are incorrect because they use different languages and syntaxes for infrastructure as code. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 559.

QUESTION 14
Which of the following technologies can be used as a central repository of Linux users and groups?

A. LDAP
B. MFA
C. SSO
D. PAM

Correct Answer: A
Section:
Explanation:
LDAP stands for Lightweight Directory Access Protocol, which is a protocol for accessing and managing a central directory of users and groups. LDAP can be used as a central repository of Linux users and groups, allowing for

centralized authentication and authorization across multiple Linux systems. MFA, SSO, and PAM are not technologies that can be used as a central repository of Linux users and groups. MFA stands for Multi-Factor Authentication, which is a method of verifying a user's identity using more than one factor, such as a password, a token, or a biometric. SSO stands for Single Sign-On, which is a feature that allows a user to log in once and access multiple applications or systems without having to re-enter credentials. PAM stands for Pluggable Authentication Modules, which is a framework that allows Linux to use different authentication methods, such as passwords, tokens, or biometrics. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing Users and Groups

QUESTION 15

A systems administrator is troubleshooting connectivity issues and trying to find out why a Linux server is not able to reach other servers on the same subnet it is connected to. When listing link parameters, the following is presented:

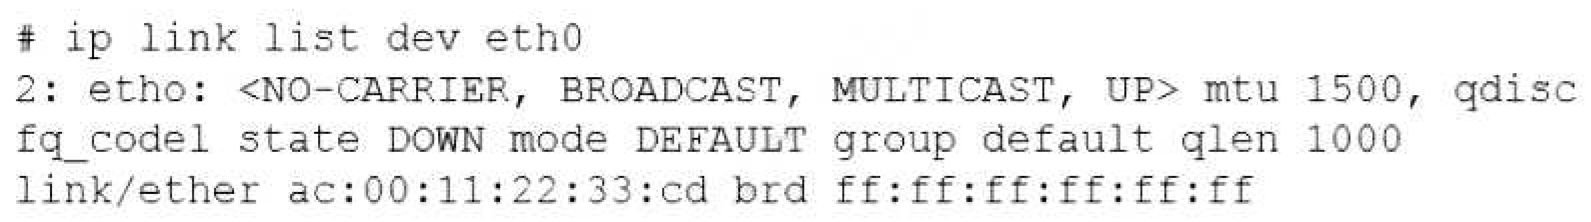

```
# ip link list dev eth0
2: etho: <NO-CARRIER, BROADCAST, MULTICAST, UP> mtu 1500, qdisc
fq_codel state DOWN mode DEFAULT group default qlen 1000
link/ether ac:00:11:22:33:cd brd ff:ff:ff:ff:ff:ff
```

Based on the output above, which of following is the MOST probable cause of the issue?

A. The address ac:00:11:22:33:cd is not a valid Ethernet address.

B. The Ethernet broadcast address should be ac:00:11:22:33:ff instead.

C. The network interface eth0 is using an old kernel module.

D. The network interface cable is not connected to a switch.

Correct Answer: D
Section:
Explanation:

The most probable cause of the connectivity issue is that the network interface cable is not connected to a switch. This can be inferred from the output of the ip link list dev eth0 command, which shows that the network interface eth0 has the NO-CARRIER flag set. This flag indicates that there is no physical link detected on the interface, meaning that the cable is either unplugged or faulty. The other options are not valid causes of the issue. The address ac:00:11:22:33:cd is a valid Ethernet address, as it follows the format of six hexadecimal octets separated by colons. The
Ethernet broadcast address should be ff:ff:ff:ff:ff:ff, which is the default value for all interfaces. The network interface eth0 is not using an old kernel module, as it shows the UP flag, which indicates that the interface is enabled and ready to transmit data. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Networking

QUESTION 16

A Linux administrator was asked to run a container with the httpd server inside. This container should be exposed at port 443 of a Linux host machine while it internally listens on port 8443. Which of the following commands will accomplish this task?

A. podman run -d -p 443:8443 httpd

B. podman run -d -p 8443:443 httpd

C. podman run -d -e 443:8443 httpd

D. podman exec -p 8443:443 httpd

Correct Answer: A
Section:
Explanation:

The command that will accomplish the task of running a container with the httpd server inside and exposing it at port 443 of the Linux host machine while it internally listens on port 8443 is podman run -d -p 443:8443 httpd. This command uses the podman tool, which is a daemonless container engine that can run and manage containers on Linux systems. The -d option runs the container in detached mode, meaning that it runs in the background without blocking the terminal. The -p option maps a port on the host machine to a port inside the container, using the format host_port:container_port. In this case, port 443 on the host machine is mapped to port 8443 inside the container, allowing external access to the httpd server. The httpd argument specifies the name of the image to run as a container, which in this case is an image that contains the Apache HTTP Server software. The other options are not correct commands for accomplishing the task. Podman run -d -p 8443:443 httpd maps port 8443 on the host machine to port 443 inside the container, which does not match the requirement. Podman run -d -e 443:8443 httpd uses the -e option instead of the -p option, which sets an environment variable inside the container instead of mapping a port. Podman exec -p 8443:443 httpd uses the podman exec command instead of the podman run command, which executes a command inside an existing container instead of creating a new one. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks

QUESTION 17
A Linux administrator needs to analyze a failing application that is running inside a container. Which of the following commands allows the Linux administrator to enter the running container and analyze the logs that are stored inside?

A. docker run -ti app /bin/sh
B. podman exec -ti app /bin/sh
C. podman run -d app /bin/bash
D. docker exec -d app /bin/bash

Correct Answer: B
Section:
Explanation:
Podman exec -ti app /bin/sh allows the Linux administrator to enter the running container and analyze the logs that are stored inside. This command uses the podman tool, which is a daemonless container engine that can run and manage containers on Linux systems. The exec option executes a command inside an existing container, in this case app, which is the name of the container that runs the failing application. The -ti option allocates a pseudo-TTY and keeps STDIN open, allowing for interactive shell access to the container. The /bin/sh argument specifies the shell command to run inside the container, which can be used to view and manipulate the log files. The other options are not correct commands for entering a running container and analyzing the logs.
Docker run -ti app /bin/sh creates a new container from the app image and runs the /bin/sh command inside it, but does not enter the existing container that runs the failing application.
Podman run -d app /bin/bash also creates a new container from the app image and runs the /bin/bash command inside it, but does so in detached mode, meaning that it runs in the background without interactive shell access. Docker exec -d app /bin/bash executes the /bin/bash command inside the existing app container, but also does so in detached mode, without interactive shell access.
Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks; View container logs | Docker Docs; How to see the logs of a docker container - Stack Overflow

QUESTION 18
A systems administrator needs to clone the partition /dev/sdc1 to /dev/sdd1. Which of the following commands will accomplish this task?

A. tar -cvzf /dev/sdd1 /dev/sdc1
B. rsync /dev/sdc1 /dev/sdd1
C. dd if=/dev/sdc1 of=/dev/sdd1
D. scp /dev/sdc1 /dev/sdd1

Correct Answer: C
Section:
Explanation:
The command dd if=/dev/sdc1 of=/dev/sdd1 copies the data from the input file (if) /dev/sdc1 to the output file (of) /dev/sdd1, byte by byte. This is the correct way to clone a partition. The other options are incorrect because they either compress the data (tar -cvzf), synchronize the files (rsync), or copy the files over a network (scp), which are not the same as cloning a partition. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 321.

QUESTION 19
When trying to log in remotely to a server, a user receives the following message:

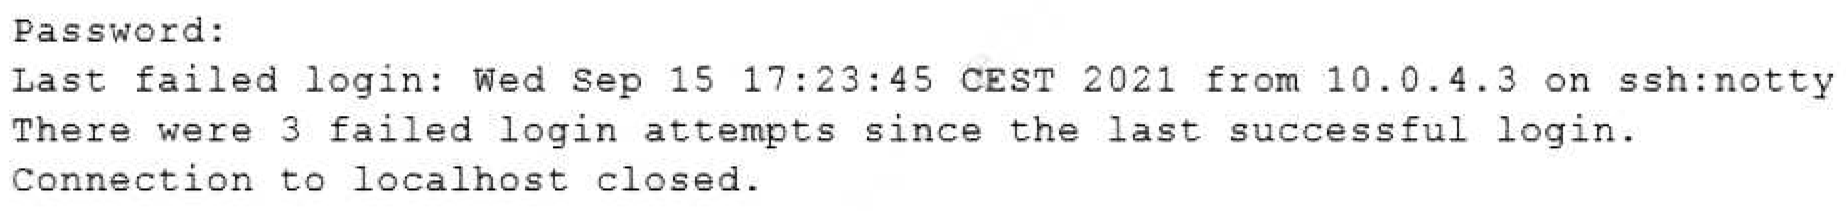

```
Password:
Last failed login: Wed Sep 15 17:23:45 CEST 2021 from 10.0.4.3 on ssh:notty
There were 3 failed login attempts since the last successful login.
Connection to localhost closed.
```

The server administrator is investigating the issue on the server and receives the following outputs:

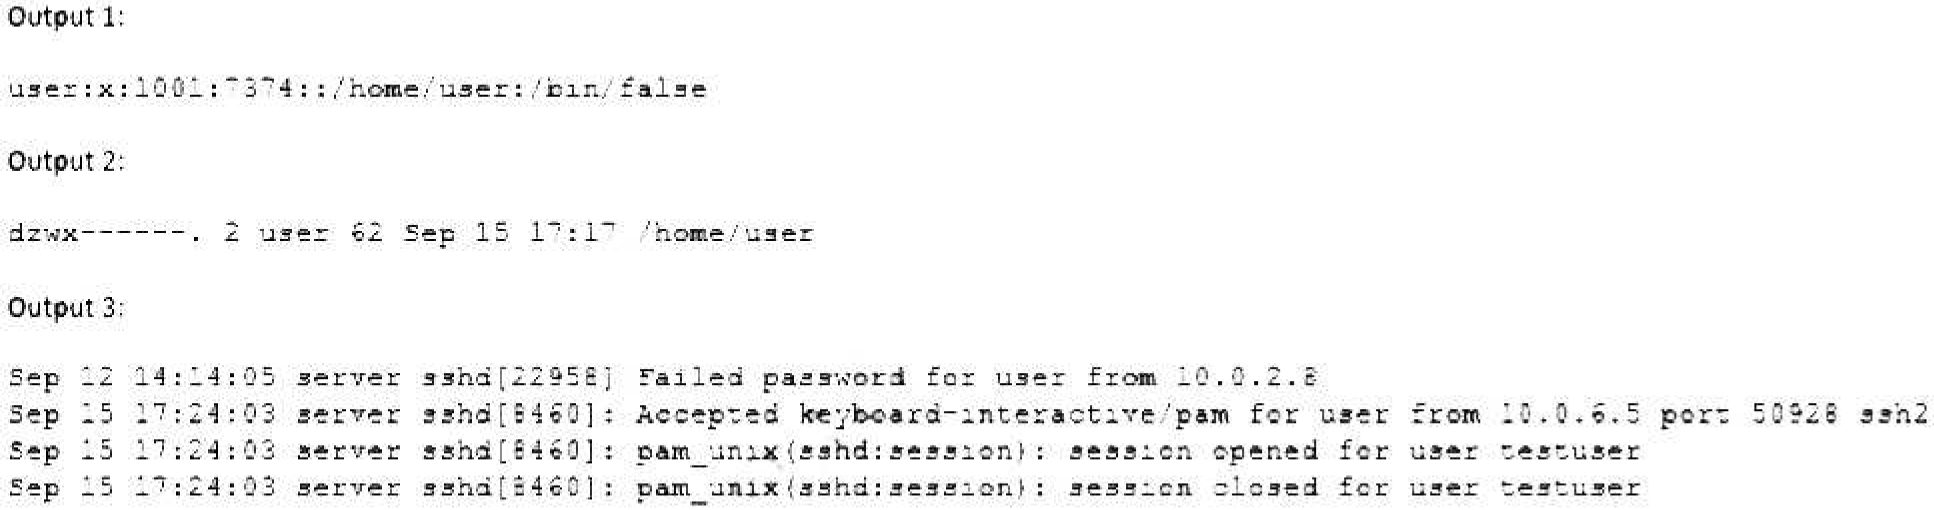

Output 1:

```
user:x:1001:7374::/home/user:/bin/false
```

Output 2:

```
dzwx------. 2 user 62 Sep 15 17:17 /home/user
```

Output 3:

```
Sep 12 14:14:05 server sshd[22958] Failed password for user from 10.0.2.8
Sep 15 17:24:03 server sshd[8460]: Accepted keyboard-interactive/pam for user from 10.0.6.5 port 50928 ssh2
Sep 15 17:24:03 server sshd[8460]: pam_unix(sshd:session): session opened for user testuser
Sep 15 17:24:03 server sshd[8460]: pam_unix(sshd:session): session closed for user testuser
```

Which of the following is causing the issue?

A. The wrong permissions are on the user's home directory.
B. The account was locked out due to three failed logins.
C. The user entered the wrong password.
D. The user has the wrong shell assigned to the account.

Correct Answer: D
Section:
Explanation:
The user has the wrong shell assigned to the account, which is causing the issue. The output 1 shows that the user's shell is set to /bin/false, which is not a valid shell and will prevent the user from logging in. The output 2 shows that the user's home directory has the correct permissions (drwxr-xrx), and the output 3 shows that the user entered the correct password and was accepted by the SSH daemon, but the session was closed immediately due to the invalid shell. The other options are incorrect because they are not supported by the outputs. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 413.

QUESTION 20
A new Linux systems administrator just generated a pair of SSH keys that should allow connection to the servers. Which of the following commands can be used to copy a key file to remote servers?
(Choose two.)

A. wget
B. ssh-keygen
C. ssh-keyscan
D. ssh-copy-id
E. ftpd
F. scp

Correct Answer: D, F
Section:
Explanation:
The commands ssh-copy-id and scp can be used to copy a key file to remote servers. The command ssh-copy-id copies the public key to the authorized_keys file on the remote server, which allows the user to log in without a password. The command scp copies files securely over SSH, which can be used to transfer the key file to any location on the remote server. The other options are incorrect because they are not related to copying key files. The command wget downloads files from the web, the command ssh-keygen generates key pairs, the command ssh-keyscan collects public keys from remote hosts, and the command ftpd is a FTP server daemon. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, pages 408-410.

QUESTION 21

A systems administrator needs to reconfigure a Linux server to allow persistent IPv4 packet forwarding. Which of the following commands is the correct way to accomplish this task?

A. echo 1 > /proc/sys/net/ipv4/ipv_forward

B. sysctl -w net.ipv4.ip_forward=1

C. firewall-cmd --enable ipv4_forwarding

D. systemctl start ipv4_forwarding

Correct Answer: B

Section:

Explanation:

The command sysctl -w net.ipv4.ip_forward=1 enables IPv4 packet forwarding temporarily by setting the kernel parameter net.ipv4.ip_forward to 1. To make this change persistent, the administrator needs to edit the file /etc/sysctl.conf and add the line net.ipv4.ip_forward = 1. The other options are incorrect because they either use the wrong file (/proc/sys/net/ipv4/ipv_forward), the wrong command (firewall-cmd or systemctl), or the wrong option (--enable or start). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 378.

QUESTION 22

A Linux administrator would like to use systemd to schedule a job to run every two hours. The administrator creates timer and service definitions and restarts the server to load these new configurations. After the restart, the administrator checks the log file and notices that the job is only running daily. Which of the following is MOST likely causing the issue?

A. The checkdiskspace.service is not running.

B. The checkdiskspace.service needs to be enabled.

C. The OnCalendar schedule is incorrect in the timer definition.

D. The system-daemon services need to be reloaded.

Correct Answer: C

Section:

Explanation:

The OnCalendar schedule is incorrect in the timer definition, which is causing the issue. The OnCalendar schedule defines when the timer should trigger the service. The format of the schedule is OnCalendar=<year>-<month>-<day> <hour>:<minute>:<second>. If any of the fields are omitted, they are assumed to be *, which means any value. Therefore, the schedule OnCalendar=*-*-* 00:00:00 means every day at midnight, which is why the job is running daily. To make the job run every two hours, the schedule should be OnCalendar=*-*-* *:00:00/2, which means every hour divisible by 2 at the start of the minute. The other options are incorrect because they are not related to the schedule. The checkdiskspace.service is running, as shown by the output of systemctl status checkdiskspace.service. The checkdiskspace.service is enabled, as shown by the output of systemctl is-enabled checkdiskspace.service. The system-daemon services do not need to be reloaded, as the timer and service definitions are already loaded by the restart. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, page 437.

QUESTION 23

An administrator deployed a Linux server that is running a web application on port 6379/tcp.
SELinux is in enforcing mode based on organization policies.
The port is open on the firewall.
Users who are trying to connect to a local instance of the web application receive Error 13, Permission denied.
The administrator ran some commands that resulted in the following output:

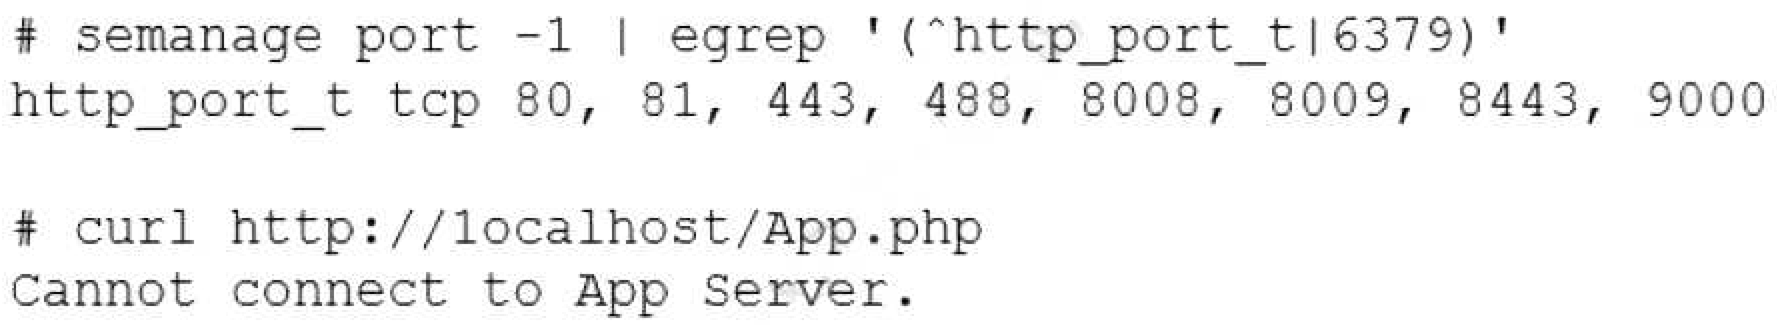

```
# semanage port -l | egrep '(^http_port_t|6379)'
http_port_t tcp 80, 81, 443, 488, 8008, 8009, 8443, 9000

# curl http://localhost/App.php
Cannot connect to App Server.
```

Which of the following commands should be used to resolve the issue?

A. semanage port -d -t http_port_t -p tcp 6379
B. semanage port -a -t http_port_t -p tcp 6379
C. semanage port -a http_port_t -p top 6379
D. semanage port -l -t http_port_tcp 6379

Correct Answer: B
Section:
Explanation:
The command semanage port -a -t http_port_t -p tcp 6379 adds a new port definition to the SELinux policy and assigns the type http_port_t to the port 6379/tcp. This allows the web application to run on this port and accept connections from users. This is the correct way to resolve the issue. The other options are incorrect because they either delete a port definition (-d), use the wrong protocol (top instead of tcp), or list the existing port definitions (-l). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 535.

QUESTION 24
A systems administrator created a web server for the company and is required to add a tag for the API so end users can connect. Which of the following would the administrator do to complete this requirement?

A. hostnamectl status --no-ask-password
B. hostnamectl set-hostname "$(perl -le "print" "A" x 86)"
C. hostnamectl set-hostname Comptia-WebNode -H root@192.168.2.14
D. hostnamectl set-hostname Comptia-WebNode --transient

Correct Answer: C
Section:
Explanation:
The command hostnamectl set-hostname Comptia-WebNode -H root@192.168.2.14 sets the hostname of the web server to Comptia-WebNode and connects to the server using the SSH protocol and the root user. This is the correct way to complete the requirement. The other options are incorrect because they either display the current hostname status (hostnamectl status), set an invalid hostname (hostnamectl set-hostname "$(perl -le "print" "A" x 86)"), or set a transient hostname that is not persistent (hostnamectl set-hostname Comptia-WebNode --transient). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 9: Managing System Components, page 291.

QUESTION 25
A systems administrator wants to back up the directory /data and all its contents to /backup/data on a remote server named remote. Which of the following commands will achieve the desired effect?

A. scp -p /data remote:/backup/data
B. ssh -i /remote:/backup/ /data
C. rsync -a /data remote:/backup/
D. cp -r /data /remote/backup/

Correct Answer: C
Section:
Explanation:
The command that will back up the directory /data and all its contents to /backup/data on a remote server named remote is rsync -a /data remote:/backup/. This command uses the rsync tool, which is a remote and local file synchronization tool. It uses an algorithm to minimize the amount of data copied by only moving the portions of files that have changed. The -a option stands for archive mode, which preserves the permissions, ownership, timestamps, and symbolic links of the files. The /data argument specifies the source directory to be backed up, and the remote:/backup/ argument specifies the destination directory on the remote server. The rsync tool will create a subdirectory named data under /backup/ on the remote server, and copy all the files and subdirectories from /data on the local server.
The other options are not correct commands for backing up a directory to a remote server. The scp -p /data remote:/backup/data command will copy the /data directory as a file named data under /backup/ on the remote server, not as a subdirectory with its contents. The -p option preserves the permissions and timestamps of the file, but not the ownership or symbolic links. The ssh -i /remote:/backup/ /data command will try to use /remote:/backup/ as an identity file for SSH authentication, which is not valid. The cp -r /data /remote/backup/ command will try to copy the /data directory to a local directory named /remote/backup/, not to a remote server. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks; rsync(1) - Linux manual page

QUESTION 26
An administrator needs to make some changes in the IaC declaration templates. Which of the following commands would maintain version control?

A.

```
git clone https://github.com/comptia/linux+-.git
git push origin
```

B.

```
git clone https://qithub.com/comptia/linux+-.git
git fetch New-Branch
```

C.

```
git clone https://github.com/comptia/linux+-.git
git status
```

D.

```
git clone https://github.com/comptia/linuxt+-.git
git checkout -b <new-branch>
```

Correct Answer: D
Section:
Explanation:
The command that will maintain version control while making some changes in the IaC declaration templates is git checkout -b <new-branch>. This command uses the git tool, which is a distributed version control system that tracks changes in source code and enables collaboration among developers. The checkout option switches to a different branch in the git repository, where a branch is a pointer to a specific commit in the history. The -b option creates a new branch with the given name, and switches to it. This way, the administrator can make changes in the new branch without affecting the main branch, and later merge them if needed.
The other options are not correct commands for maintaining version control while making some changes in the IaC declaration templates. The git clone https://github.com/comptia/linux±.git command will clone an existing repository from a remote URL to a local directory, but it will not create a new branch for making changes. The git push origin command will push the local changes to a remote repository named origin, but it will not create a new branch for making changes. The git fetch New-Branch command will fetch updates from a remote branch named New-Branch, but it will not create a new branch for making changes. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Source Code; Git - Basic Branching and Merging

QUESTION 27
After listing the properties of a system account, a systems administrator wants to remove the expiration date of a user account. Which of the following commands will accomplish this task?

A. chgrp system accountname
B. passwd -s accountname
C. chmod -G system account name
D. chage -E -1 accountname

Correct Answer: D
Section:

Explanation:
The command chage -E -1 accountname will accomplish the task of removing the expiration date of a user account. The chage command is a tool for changing user password aging information on Linux systems. The -E option sets the expiration date of the user account, and the -1 value means that the account will never expire. The command chage -E -1 accountname will remove the expiration date of the user account named accountname. This is the correct command to use to accomplish the task.
The other options are incorrect because they either do not affect the expiration date (chgrp, passwd, or chmod) or do not exist (chmod -G). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Users and Groups, page 467.

QUESTION 28
A systems administrator wants to be sure the sudo rules just added to /etc/sudoers are valid. Which of the following commands can be used for this task?

A. visudo -c
B. test -f /etc/sudoers
C. sudo vi check
D. cat /etc/sudoers | tee test

Correct Answer: A
Section:
Explanation:
The command visudo -c can be used to check the validity of the sudo rules in the /etc/sudoers file.
The visudo command is a tool for editing and validating the /etc/sudoers file, which defines the rules for the sudo command. The -c option checks the syntax and logic of the file and reports any errors or warnings. The command visudo -c will verify the sudo rules and help the administrator avoid any mistakes. This is the correct command to use for this task. The other options are incorrect because they either do not check the validity of the file (test, sudo, or cat) or do not exist (sudo vi check). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 546.

QUESTION 29
A user generated a pair of private-public keys on a workstation. Which of the following commands will allow the user to upload the public key to a remote server and enable passwordless login?

A. scp ~/.ssh/id_rsa user@server:~/
B. rsync ~ /.ssh/ user@server:~/
C. ssh-add user server
D. ssh-copy-id user@server

Correct Answer: D
Section:
Explanation:
The command ssh-copy-id user@server will allow the user to upload the public key to a remote server and enable passwordless login. The ssh-copy-id command is a tool for copying the public key to a remote server and appending it to the authorized_keys file, which is used for public key authentication. The command will also set the appropriate permissions on the remote server to ensure the security of the key. The command ssh-copy-id user@server will copy the public key of the user to the server and allow the user to log in without a password. This is the correct command to use for this task. The other options are incorrect because they either do not copy the public key (scp, rsync, or ssh-add) or do not use the correct syntax (scp ~/.ssh/id_rsa user@server:~/ instead of scp ~/.ssh/id_rsa.pub user@server:~/ or rsync ~ /.ssh/ user@server:~/ instead of rsync ~/.ssh/id_rsa.pub user@server:~/). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 410.

QUESTION 30
A Linux administrator created a new file system. Which of the following files must be updated to ensure the filesystem mounts at boot time?

A. /etc/sysctl
B. /etc/filesystems
C. /etc/fstab
D. /etc/nfsmount.conf

Correct Answer: C
Section:
Explanation:
The file that must be updated to ensure the filesystem mounts at boot time is /etc/fstab. This file contains information about the filesystems that are mounted automatically by the mount -a command, which is usually invoked during the system startup. The /etc/fstab file has six fields for each filesystem: device name, mount point, filesystem type, mount options, dump frequency, and pass number. To add a new filesystem to the /etc/fstab file, you need to specify these fields correctly and make sure the mount point directory exists.
The other options are not correct files for controlling persistent mount points of filesystems. The /etc/sysctl file is used to configure kernel parameters at runtime. The /etc/filesystems file is used to specify the order of filesystem types used by mount when no filesystem type is given. The /etc/nfsmount.conf file is used to set options for mounting NFS filesystems. Reference: Persistently mounting file systems; fstab(5) - Linux manual page

QUESTION 31
A Linux administrator is troubleshooting a memory-related issue. Based on the output of the commands:

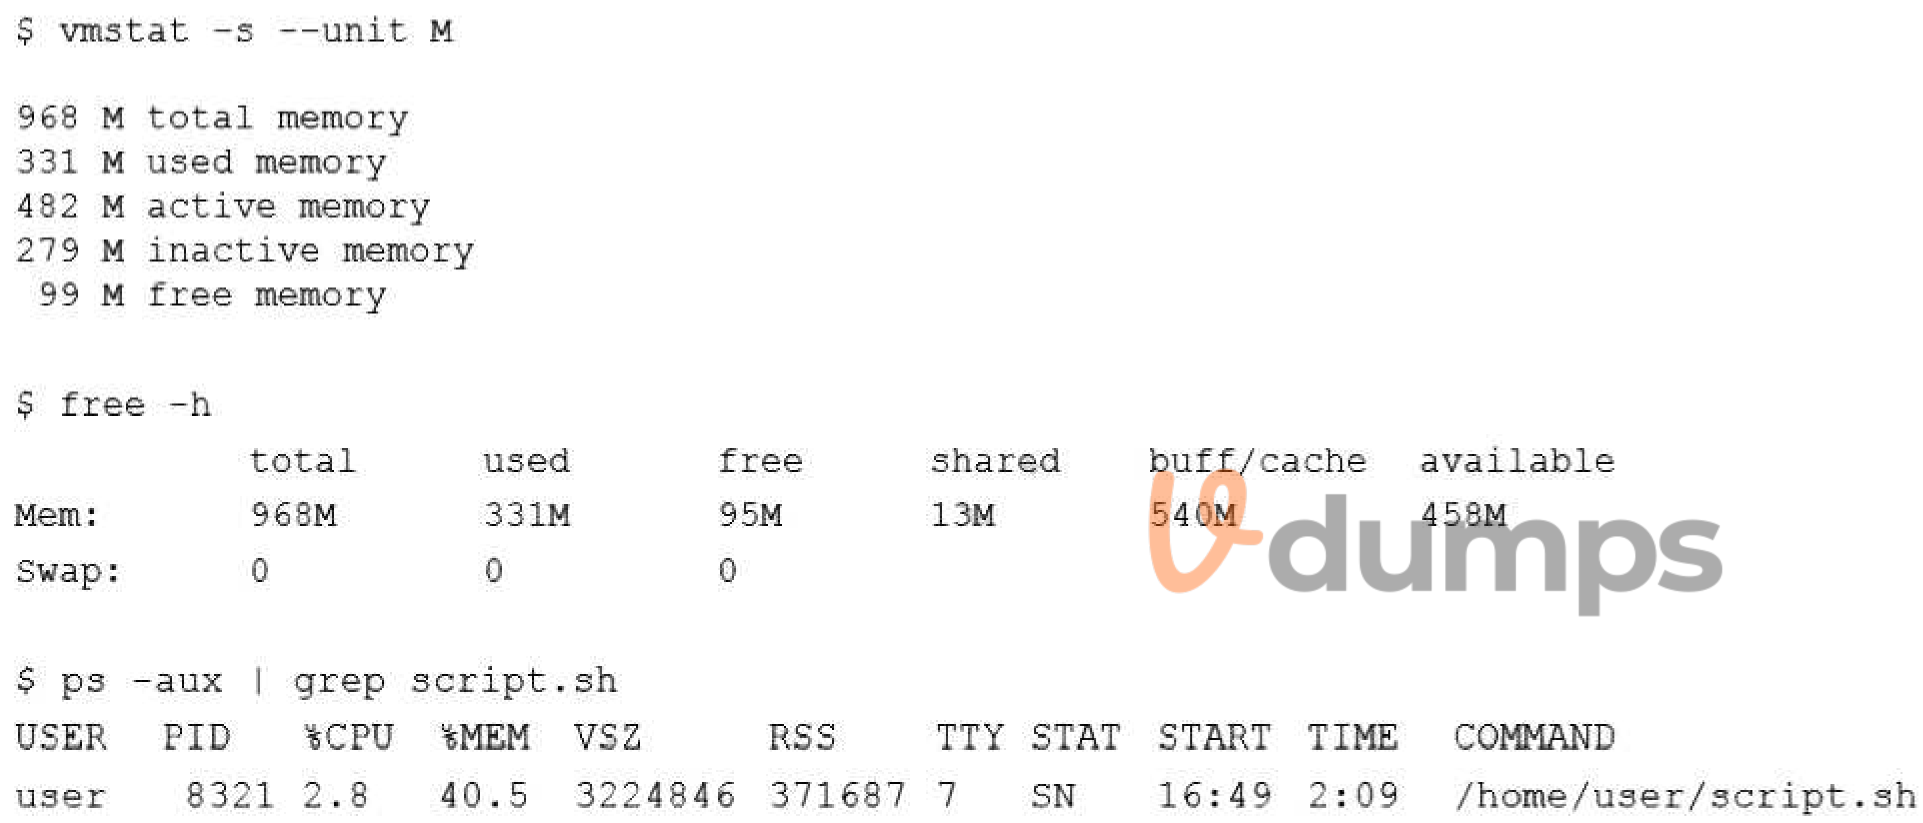

```
$ vmstat -s --unit M

968 M total memory
331 M used memory
482 M active memory
279 M inactive memory
 99 M free memory


$ free -h
          total     used      free     shared    buff/cache  available
Mem:      968M      331M      95M      13M       540M        458M
Swap:     0         0         0


$ ps -aux | grep script.sh
USER   PID   %CPU  %MEM  VSZ      RSS     TTY  STAT  START  TIME   COMMAND
user    8321 2.8   40.5  3224846  371687  7    SN    16:49  2:09   /home/user/script.sh
```

Which of the following commands would address the issue?

A. top -p 8321
B. kill -9 8321
C. renice -10 8321
D. free 8321

Correct Answer: B
Section:
Explanation:
The command that would address the memory-related issue is kill -9 8321. This command will send a SIGKILL signal to the process with the PID 8321, which is the mysqld process that is using 99.7% of the available memory according to the top output. The SIGKILL signal will terminate the process immediately and free up the memory it was using. However, this command should be used with caution as it may cause data loss or corruption if the process was performing some critical operations.
The other options are not correct commands for addressing the memory-related issue. The top -p 8321 command will only display information about the process with the PID 8321, but will not kill it or reduce its memory usage. The renice -10 8321 command will change the priority (niceness) of the process with the PID 8321 to -10, which means it will have a higher scheduling priority, but this will not affect its memory consumption. The free 8321 command is invalid because free does not take a PID as an argument; free only displays information about the total, used, and free memory in the system. Reference: How to troubleshoot Linux server memory issues;

kill(1) - Linux manual page

QUESTION 32
A systems administrator made some unapproved changes prior to leaving the company. The newly hired administrator has been tasked with revealing the system to a compliant state. Which of the following commands will list and remove the correspondent packages?

A. dnf list and dnf remove last
B. dnf remove and dnf check
C. dnf info and dnf upgrade
D. dnf history and dnf history undo last

Correct Answer: D
Section:
Explanation:
The commands that will list and remove the corresponding packages are dnf history and dnf history undo last. The dnf history command will display a list of all transactions performed by dnf, such as installing, updating, or removing packages. Each transaction has a unique ID, a date and time, an action, and a number of altered packages. The dnf history undo last command will undo the last transaction performed by dnf, meaning that it will reverse all package changes made by that transaction. For example, if the last transaction installed some packages, dnf history undo last will remove them.
The other options are not correct commands for listing and removing corresponding packages. The dnf list command will display a list of available packages in enabled repositories, but not the packages installed by dnf transactions. The dnf remove command will remove specified packages from the system, but not all packages from a specific transaction. The dnf info command will display detailed information about specified packages, but not about dnf transactions. The dnf upgrade command will upgrade all installed packages to their latest versions, but not undo any package changes. Reference: Handling package management history; dnf(8) - Linux manual page

QUESTION 33
An administrator transferred a key for SSH authentication to a home directory on a remote server.
The key file was moved to .ssh/authorized_keys location in order to establish SSH connection without a password. However, the SSH command still asked for the password. Given the following output:

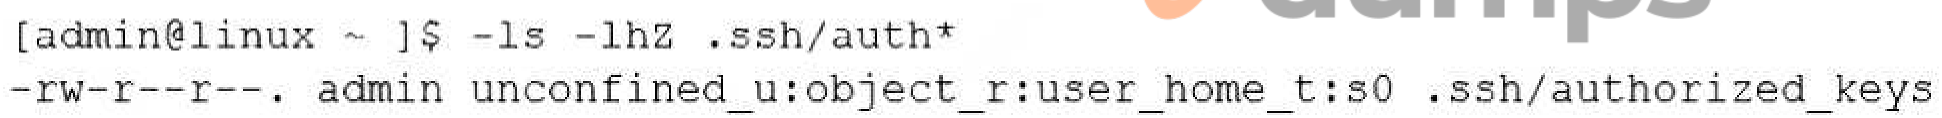

```
[admin@linux ~ ]$ -ls -lhZ .ssh/auth*
-rw-r--r--. admin unconfined_u:object_r:user_home_t:s0 .ssh/authorized_keys
```

Which of the following commands would resolve the issue?

A. restorecon .ssh/authorized_keys
B. ssh_keygen -t rsa -o .ssh/authorized_keys
C. chown root:root .ssh/authorized_keys
D. chmod 600 .ssh/authorized_keys

Correct Answer: D
Section:
Explanation:
The command that would resolve the issue is chmod 600 .ssh/authorized_keys. This command will change the permissions of the .ssh/authorized_keys file to 600, which means that only the owner of the file can read and write it. This is necessary for SSH key authentication to work properly, as SSH will refuse to use a key file that is accessible by other users or groups for security reasons. The output of ls -l shows that currently the .ssh/authorized_keys file has permissions of 664, which means that both the owner and group can read and write it, and others can read it.
The other options are not correct commands for resolving the issue. The restorecon .ssh/authorized_keys command will restore the default SELinux security context for the .ssh/authorized_keys file, but this will not change its permissions or ownership. The ssh_keygen -t rsa -o .ssh/authorized_keys command is invalid because ssh_keygen is not a valid command (the correct command is ssh-keygen), and the -o option is used to specify a new output format for the key file, not the output file name. The chown root:root .ssh/authorized_keys command will change the owner and group of the .ssh/authorized_keys file to root, but this will not change its permissions or make it accessible by the user who wants to log in with SSH key authentication. Reference: How to Use Public Key Authentication with SSH; chmod(1) - Linux manual page

QUESTION 34

A cloud engineer needs to remove all dangling images and delete all the images that do not have an associated container. Which of the following commands will help to accomplish this task?

A. docker images prune -a
B. docker push images -a
C. docker rmi -a images
D. docker images rmi --all

Correct Answer: A
Section:
Explanation:
The command docker images prune -a will help to remove all dangling images and delete all the images that do not have an associated container. The docker command is a tool for managing Docker containers and images. The images subcommand operates on images. The prune option removes unused images. The -a option removes all images, not just dangling ones. A dangling image is an image that is not tagged and is not referenced by any container. This command will accomplish the task of cleaning up the unused images. The other options are incorrect because they either do not exist (docker push images -a or docker images rmi --all) or do not remove images (docker rmi -a images only removes images that match the name or ID of "images"). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 567.

QUESTION 35
A Linux system is failing to boot with the following error:

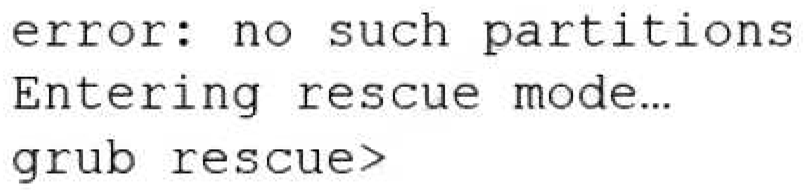

```
error: no such partitions
Entering rescue mode…
grub rescue>
```

Which of the following actions will resolve this issue? (Choose two.)

A. Execute grub-install --root-directory=/mnt and reboot.
B. Execute grub-install /dev/sdX and reboot.
C. Interrupt the boot process in the GRUB menu and add rescue to the kernel line.
D. Fix the partition modifying /etc/default/grub and reboot.
E. Interrupt the boot process in the GRUB menu and add single to the kernel line.
F. Boot the system on a LiveCD/ISO.

Correct Answer: B, F
Section:
Explanation:
The administrator should do the following two actions to resolve the issue:
Boot the system on a LiveCD/ISO. This is necessary to access the system and repair the boot loader. A LiveCD/ISO is a bootable media that contains a Linux distribution that can run without installation. The administrator can boot the system from the LiveCD/ISO and mount the root partition of the system to a temporary directory, such as /mnt.
Execute grub-install /dev/sdX and reboot. This will reinstall the GRUB boot loader to the disk device, where sdX is the device name of the disk, such as sda or sdb. The GRUB boot loader is a program that runs when the system is powered on and allows the user to choose which operating system or kernel to boot. The issue is caused by a corrupted or missing GRUB boot loader, which prevents the system from booting. The command grub-install will restore the GRUB boot loader and fix the issue.
The other options are incorrect because they either do not fix the boot loader (interrupt the boot process in the GRUB menu or fix the partition modifying /etc/default/grub) or do not use the correct syntax (grub-install --root-directory=/mnt instead of grub-install /dev/sdX or rescue or single instead of recovery in the GRUB menu). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing the Linux Boot Process, pages 265-266.

QUESTION 36
A Linux administrator needs to create an image named sda.img from the sda disk and store it in the /tmp directory. Which of the following commands should be used to accomplish this task?

A. dd of=/dev/sda if=/tmp/sda.img

B. dd if=/dev/sda of=/tmp/sda.img

C. dd --if=/dev/sda --of=/tmp/sda.img

D. dd --of=/dev/sda --if=/tmp/sda.img

Correct Answer: B
Section:
Explanation:
The command dd if=/dev/sda of=/tmp/sda.img should be used to create an image named sda.img from the sda disk and store it in the /tmp directory. The dd command is a tool for copying and converting data on Linux systems. The if option specifies the input file or device, in this case /dev/sda, which is the disk device. The of option specifies the output file or device, in this case /tmp/sda.img, which is the image file. The command dd if=/dev/sda of=/tmp/sda.img will copy the entire disk data from /dev/sda to /tmp/sda.img and create an image file. This is the correct command to use to accomplish the task. The other options are incorrect because they either use the wrong options (--if or --of instead of if or of) or swap the input and output (dd of=/dev/sda if=/tmp/sda.img or dd --of=/dev/sda --if=/tmp/sda.img). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 323.

QUESTION 37

A Linux administrator is creating a primary partition on the replacement hard drive for an application server. Which of the following commands should the administrator issue to verify the device name of this partition?

A. sudo fdisk /dev/sda

B. sudo fdisk -s /dev/sda

C. sudo fdisk -l

D. sudo fdisk -h

Correct Answer: C
Section:
Explanation:
The command sudo fdisk -l should be issued to verify the device name of the partition.
The sudo command allows the administrator to run commands as the superuser or another user.
The fdisk command is a tool for manipulating disk partitions on Linux systems. The -l option lists the partitions on all disks or a specific disk. The command sudo fdisk -l will show the device names, sizes, types, and other information of the partitions on all disks. The administrator can identify the device name of the partition by looking at the output. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not list the partitions (sudo fdisk /dev/sda or sudo fdisk -h) or do not exist (sudo fdisk -s /dev/sda). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 317.

QUESTION 38

A systems administrator is investigating why one of the servers has stopped connecting to the internet.

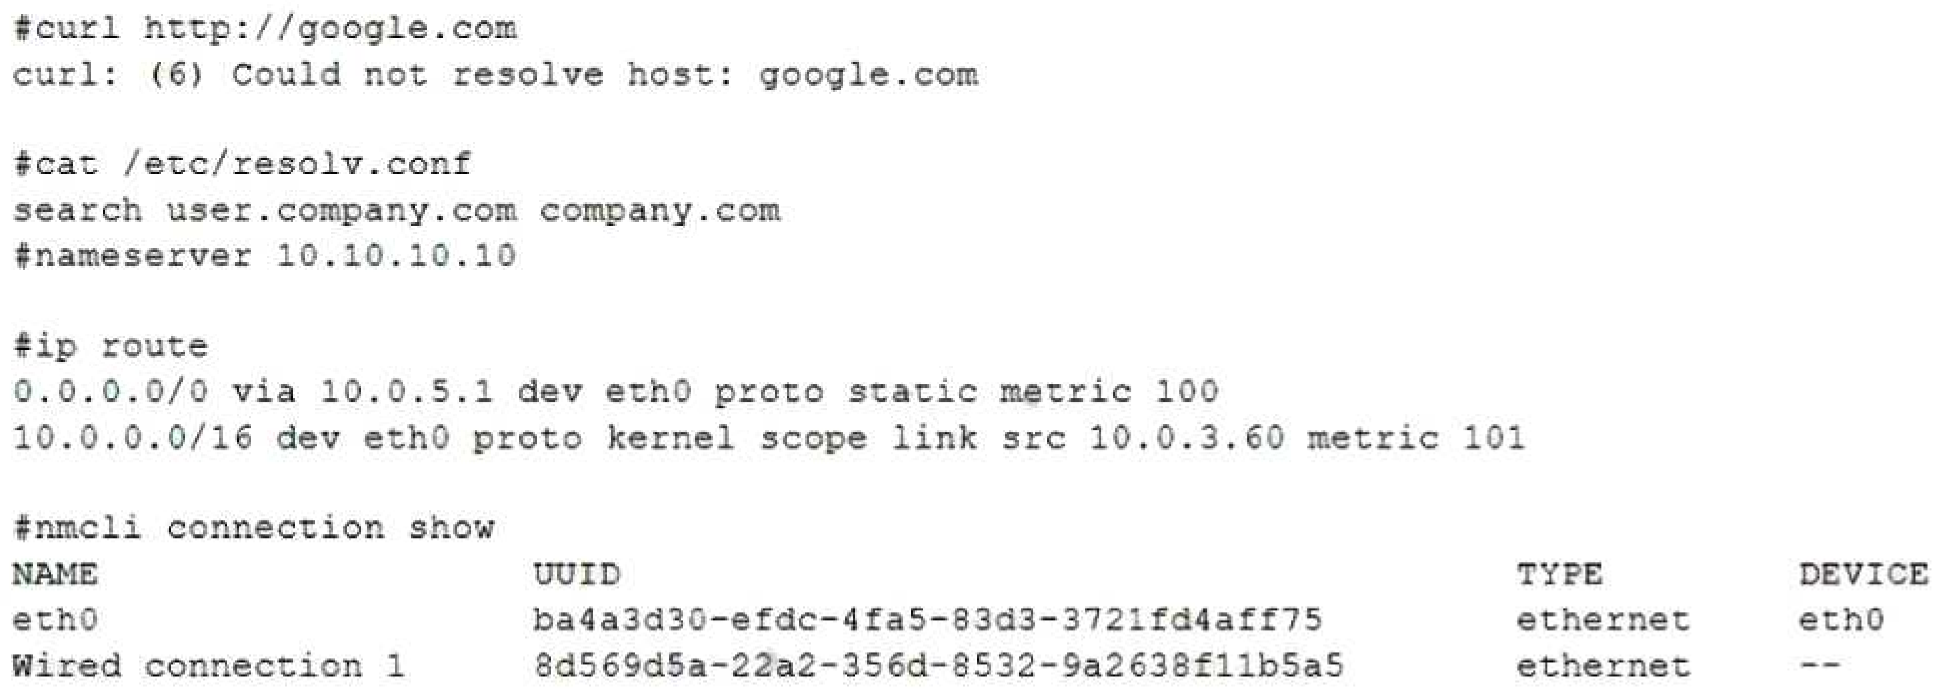

```
#curl http://google.com
curl: (6) Could not resolve host: google.com

#cat /etc/resolv.conf
search user.company.com company.com
#nameserver 10.10.10.10

#ip route
0.0.0.0/0 via 10.0.5.1 dev eth0 proto static metric 100
10.0.0.0/16 dev eth0 proto kernel scope link src 10.0.3.60 metric 101

#nmcli connection show
NAME                   UUID                                    TYPE        DEVICE
eth0                   ba4a3d30-efdc-4fa5-83d3-3721fd4aff75    ethernet    eth0
Wired connection 1     8d569d5a-22a2-356d-8532-9a2638f11b5a5    ethernet    --
```

Which of the following is causing the issue?

A. The DNS address has been commented out in the configuration file.

B. The search entry in the /etc/resolv.conf file is incorrect.

C. Wired connection 1 is offline.

D. No default route is defined.

Correct Answer: D
Section:
Explanation:
The issue is caused by the lack of a default route defined in the /etc/sysconfig/network-scripts/ifcfgenp0s3 file. A default route is a special route that specifies where to send packets that do not match any other routes in the routing table. Without a default route, the server will not be able to communicate with hosts outside its local network. The default route is usually configured with the GATEWAY option in the network interface configuration file. For example, to set the default gateway to 192.168.1.1, the file should contain:
GATEWAY=192.168.1.1
The other options are not causing the issue. The DNS address is not commented out in the configuration file, it is specified with the DNS1 option. The search entry in the /etc/resolv.conf file is correct, it specifies the domain name to append to unqualified hostnames. Wired connection 1 is online, as indicated by the ONBOOT=yes option and the output of ip link show enp0s3 command. Reference: Configuring IP Networking with nmcli; Configuring IP Networking with ifcfg Files

QUESTION 39
A systems administrator is tasked with installing GRUB on the legacy MBR of the SATA hard drive.
Which of the following commands will help the administrator accomplish this task?

A. grub-install /dev/hda

B. grub-install /dev/sda

C. grub-install /dev/sr0

D. grub-install /dev/hd0,0

Correct Answer: B
Section:
Explanation:

The command that will help the administrator install GRUB on the legacy MBR of the SATA hard drive is grub-install /dev/sda. This command will install GRUB on the master boot record (MBR) of the first SATA disk (/dev/sda). The MBR is the first sector of a disk that contains boot code and a partition table. GRUB will overwrite the boot code and place its own code that can load GRUB modules and configuration files from a specific partition.
The other options are not correct commands for installing GRUB on the legacy MBR of the SATA hard drive. The grub-install /dev/hda command will try to install GRUB on the first IDE disk (/dev/hda), which may not exist or may not be bootable. The grub-install /dev/sr0 command will try to install
GRUB on the first SCSI CD-ROM device (/dev/sr0), which is not a hard drive and may not be bootable.
The grub-install /dev/hd0,0 command is invalid because grub-install does not accept partition names as arguments, only disk names. Reference: Installing GRUB using grub-install; GRUB Manual

QUESTION 40
A junior Linux administrator is tasked with installing an application. The installation guide states the application should only be installed in a run level 5 environment.

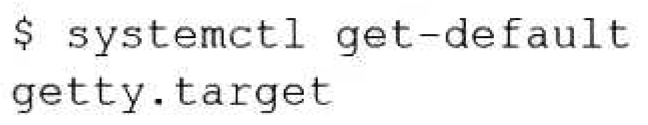

```
$ systemctl get-default
getty.target
```

Which of the following commands would ensure the server is set to runlevel 5?

A. systemctl isolate multi-user.target
B. systemctl isolate graphical.target
C. systemctl isolate network.target
D. systemctl isolate basic.target

Correct Answer: B
Section:
Explanation:
The command that would ensure the server is set to runlevel 5 is systemctl isolate graphical.target. This command will change the current target (or runlevel) of systemd to graphical.target, which is equivalent to runlevel 5 in SysV init systems. Graphical.target means that the system will start with a graphical user interface (GUI) and all services required for it.
The other options are not correct commands for setting the server to runlevel 5. The systemctl isolate multi-user.target command will change the current target to multi-user.target, which is equivalent to runlevel 3 in SysV init systems. Multi-user.target means that the system will start with multiple user logins and networking, but without a GUI. The systemctl isolate network.target command will change the current target to network.target, which is not a real runlevel but a synchronization point for network-related services. Network.target means that network functionality should be available, but does not specify whether it should be started before or after it. The systemctl isolate basic.target command will change the current target to basic.target, which is also not a real runlevel but a synchronization point for basic system services. Basic.target means that all essential services should be started, but does not specify whether it should be started before or after it. Reference: systemd System and Service Manager; systemd.special(7) - Linux manual page

QUESTION 41
A Linux administrator is tasked with adding users to the system. However, the administrator wants to ensure the users' access will be disabled once the project is over. The expiration date should be 2021-09-30. Which of the following commands will accomplish this task?

A. sudo useradd -e 2021-09-30 Project_user
B. sudo useradd -c 2021-09-30 Project_user
C. sudo modinfo -F 2021-09-30 Project_uses
D. sudo useradd -m -d 2021-09-30 Project_user

Correct Answer: A
Section:
Explanation:
The command that will accomplish this task is sudo useradd -e 2021-09-30 Project_user. This command will create a new user account named Project_user with an expiration date of 2021-09-30. The -e option of useradd specifies the date on which the user account will be disabled in YYYY-MMDD format.
The other options are not correct commands for creating a user account with an expiration date. The sudo useradd -c 2021-09-30 Project_user command will create a new user account named Project_user with a comment of 2021-09-30. The -c option of useradd specifies a comment or description for the user account, not an expiration date. The sudo modinfo -F 2021-09-30 Project_user command is invalid because modinfo is not a command for managing user accounts, but a command for displaying information about kernel modules. The -F option of modinfo specifies a field name to show, not an expiration date. The sudo useradd -m -d 2021-09-30 Project_user command will create a new user account named Project_user with a home directory of 2021-09-30.

The -m option of useradd specifies that the home directory should be created if it does not exist, and the -d option specifies the home directory name, not an expiration date. Reference: useradd(8) -Linux manual page; modinfo(8) - Linux manual page

QUESTION 42
A DevOps engineer needs to download a Git repository from
https://git.company.com/admin/project.git. Which of the following commands will achieve this goal?

A. git clone https://git.company.com/admin/project.git
B. git checkout https://git.company.com/admin/project.git
C. git pull https://git.company.com/admin/project.git
D. git branch https://git.company.com/admin/project.git

Correct Answer: A
Section:
Explanation:
The command git clone https://git.company.com/admin/project.git will achieve the goal of downloading a Git repository from the given URL. The git command is a tool for managing version control systems. The clone option creates a copy of an existing repository. The URL specifies the location of the repository to clone, in this case https://git.company.com/admin/project.git. The command git clone https://git.company.com/admin/project.git will download the repository and create a directory named project in the current working directory. This is the correct command to use to accomplish the goal. The other options are incorrect because they either do not download the repository (git checkout, git pull, or git branch) or do not use the correct syntax (git checkout
https://git.company.com/admin/project.git instead of git checkout -b project
https://git.company.com/admin/project.git or git branch
https://git.company.com/admin/project.git instead of git branch project
https://git.company.com/admin/project.git). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 571.

QUESTION 43
An administrator installed an application from source into /opt/operations1/ and has received numerous reports that users are not able to access the application without having to use the full path /opt/operations1/bin/*. Which of the following commands should be used to resolve this issue?

A. echo 'export PATH=$PATH:/opt/operations1/bin' >> /etc/profile
B. echo 'export PATH=/opt/operations1/bin' >> /etc/profile
C. echo 'export PATH=$PATH/opt/operations1/bin' >> /etc/profile
D. echo 'export $PATH:/opt/operations1/bin' >> /etc/profile

Correct Answer: A
Section:
Explanation:
The command echo 'export PATH=$PATH:/opt/operations1/bin' >> /etc/profile should be used to resolve the issue of users not being able to access the application without using the full path.
The echo command prints the given string to the standard output. The export command sets an environment variable and makes it available to all child processes. The PATH variable contains a list of directories where the shell looks for executable files. The $PATH expands to the current value of the PATH variable. The : separates the directories in the list. The /opt/operations1/bin is the directory where the application is installed. The >> operator appends the output to the end of the file.
The /etc/profile file is a configuration file that is executed when a user logs in. The command echo 'export PATH=$PATH:/opt/operations1/bin' >> /etc/profile will add the /opt/operations1/bin directory to the PATH variable for all users and allow them to access the application without using the full path. This is the correct command to use to resolve the issue. The other options are incorrect because they either overwrite the PATH variable (echo 'export PATH=/opt/operations1/bin' >> /etc/profile) or do not use the correct syntax (echo 'export PATH=$PATH/opt/operations1/bin' >> /etc/profile or echo 'export $PATH:/opt/operations1/bin' >> /etc/profile). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 9: Working with the Linux Shell, page 295.

QUESTION 44
A Linux system is getting an error indicating the root filesystem is full. Which of the following commands should be used by the systems administrator to resolve this issue? (Choose three.)

A. df -h /

B. fdisk -1 /dev/sdb
C. growpart /dev/mapper/rootvg-rootlv
D. pvcreate /dev/sdb
E. lvresize -L +10G -r /dev/mapper/rootvg-rootlv
F. lsblk /dev/sda
G. parted -l /dev/mapper/rootvg-rootlv
H. vgextend /dev/rootvg /dev/sdb

Correct Answer: A, C, E
Section:
Explanation:
The administrator should use the following three commands to resolve the issue of the root filesystem being full:
df -h /. This command will show the disk usage of the root filesystem in a human-readable format. The df command is a tool for reporting file system disk space usage. The -h option displays the sizes in powers of 1024 (e.g., 1K, 234M, 2G). The / specifies the root filesystem. The command df -h / will show the total size, used space, available space, and percentage of the root filesystem. This command will help the administrator identify the problem and plan the solution.
growpart /dev/mapper/rootvg-rootlv. This command will grow the partition that contains the root filesystem to the maximum size available. The growpart command is a tool for resizing partitions on Linux systems. The /dev/mapper/rootvg-rootlv is the device name of the partition, which is a logical volume managed by the Logical Volume Manager (LVM). The command growpart /dev/mapper/rootvg-rootlv will extend the partition to fill the disk space and increase the size of the root filesystem. This command will help the administrator solve the problem and free up space.
lvresize -L +10G -r /dev/mapper/rootvg-rootlv. This command will resize the logical volume that contains the root filesystem and add 10 GB of space. The lvresize command is a tool for resizing logical volumes on Linux systems. The -L option specifies the new size of the logical volume, in this case +10G, which means 10 GB more than the current size. The -r option resizes the underlying file system as well. The /dev/mapper/rootvg-rootlv is the device name of the logical volume, which is the same as the partition name. The command lvresize -L +10G -r /dev/mapper/rootvg-rootlv will increase the size of the logical volume and the root filesystem by 10 GB and free up space. This command will help the administrator solve the problem and free up space.
The other options are incorrect because they either do not affect the root filesystem (fdisk -1 /dev/sdb, pvcreate /dev/sdb, lsblk /dev/sda, or vgextend /dev/rootvg /dev/sdb) or do not use the correct syntax (fdisk -1 /dev/sdb instead of fdisk -l /dev/sdb or parted -l /dev/mapper/rootvgrootlv instead of parted /dev/mapper/rootvg-rootlv print). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 318-319, 331-332.

QUESTION 45
A cloud engineer is asked to copy the file deployment.yaml from a container to the host where the container is running. Which of the following commands can accomplish this task?

A. docker cp container_id/deployment.yaml deployment.yaml
B. docker cp container_id:/deployment.yaml deployment.yaml
C. docker cp deployment.yaml local://deployment.yaml
D. docker cp container_id/deployment.yaml local://deployment.yaml

Correct Answer: B
Section:
Explanation:
The command docker cp container_id:/deployment.yaml deployment.yaml can accomplish the task of copying the file deployment.yaml from a container to the host. The docker command is a tool for managing Docker containers and images. The cp option copies files or directories between a container and the local filesystem. The container_id is the identifier of the container, which can be obtained by using the docker ps command. The /deployment.yaml is the path of the file in the container, which must be preceded by a slash. The deployment.yaml is the path of the file on the host, which can be relative or absolute. The command docker cp container_id:/deployment.yaml deployment.yaml will copy the file deployment.yaml from the container to the current working directory on the host. This is the correct command to use to accomplish the task. The other options are incorrect because they either use the wrong syntax (docker cp container_id/deployment.yaml deployment.yaml or docker cp container_id/deployment.yaml local://deployment.yaml) or do not exist (docker cp deployment.yaml local://deployment.yaml). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 567.

QUESTION 46
A Linux system is failing to start due to issues with several critical system processes. Which of the following options can be used to boot the system into the single user mode? (Choose two.)

A. Execute the following command from the GRUB rescue shell: mount -o remount, ro/sysroot.

B. Interrupt the boot process in the GRUB menu and add systemd.unit=single in the kernel line.
C. Interrupt the boot process in the GRUB menu and add systemd.unit=rescue.target in the kernel line.
D. Interrupt the boot process in the GRUB menu and add single=user in the kernel line.
E. Interrupt the boot process in the GRUB menu and add init=/bin/bash in the kernel line.
F. Interrupt the boot process in the GRUB menu and add systemd.unit=single.target in the kernel line.

Correct Answer: C, F
Section:
Explanation:
The administrator can use the following two options to boot the system into the single user mode:
Interrupt the boot process in the GRUB menu and add systemd.unit=rescue.target in the kernel line.
This option will boot the system into the rescue mode, which is a minimal environment that allows the administrator to perform basic tasks such as repairing the system. The GRUB menu is a screen that appears when the system is powered on and allows the administrator to choose which kernel or operating system to boot. The kernel line is a line that specifies the parameters for the kernel, such as the root device, the init system, and the boot options. The administrator can interrupt the boot process by pressing the e key in the GRUB menu and edit the kernel line by adding systemd.unit=rescue.target at the end. This option will tell the system to use the rescue target, which is a unit that defines the state of the system in the rescue mode. The administrator can then press Ctrl+X to boot the system with the modified kernel line. This option will boot the system into the single user mode and allow the administrator to troubleshoot the issues.
Interrupt the boot process in the GRUB menu and add systemd.unit=single.target in the kernel line.
This option will boot the system into the single user mode, which is a mode that allows the administrator to log in as the root user and perform maintenance tasks. The GRUB menu and the kernel line are the same as the previous option. The administrator can interrupt the boot process by pressing the e key in the GRUB menu and edit the kernel line by adding systemd.unit=single.target at the end. This option will tell the system to use the single target, which is a unit that defines the state of the system in the single user mode. The administrator can then press Ctrl+X to boot the system with the modified kernel line. This option will boot the system into the single user mode and allow the administrator to troubleshoot the issues.
The other options are incorrect because they either do not boot the system into the single user mode (execute the following command from the GRUB rescue shell: mount -o remount, ro/sysroot or interrupt the boot process in the GRUB menu and add systemd.unit=single in the kernel line) or do not use the correct syntax (interrupt the boot process in the GRUB menu and add single=user in the kernel line or interrupt the boot process in the GRUB menu and add init=/bin/bash in the kernel line). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8:
Managing the Linux Boot Process, pages 267-268.

QUESTION 47
A DevOps engineer needs to allow incoming traffic to ports in the range of 4000 to 5000 on a Linux server. Which of the following commands will enforce this rule?

A. iptables -f filter -I INPUT -p tcp --dport 4000:5000 -A ACCEPT
B. iptables -t filter -A INPUT -p tcp --dport 4000:5000 -j ACCEPT
C. iptables filter -A INPUT -p tcp --dport 4000:5000 -D ACCEPT
D. iptables filter -S INPUT -p tcp --dport 4000:5000 -A ACCEPT

Correct Answer: B
Section:
Explanation:
The command iptables -t filter -A INPUT -p tcp --dport 4000:5000 -j ACCEPT will enforce the rule of allowing incoming traffic to ports in the range of 4000 to 5000 on a Linux server.
The iptables command is a tool for managing firewall rules on Linux systems. The -t option specifies the table to operate on, in this case filter, which is the default table that contains the rules for filtering packets. The -A option appends a new rule to the end of a chain, in this case INPUT, which is the chain that processes the packets that are destined for the local system. The -p option specifies the protocol to match, in this case tcp, which is the transmission control protocol. The --dport option specifies the destination port or port range to match, in this case 4000:5000, which is the range of ports from 4000 to 5000. The -j option specifies the target to jump to if the rule matches, in this case ACCEPT, which is the target that allows the packet to pass through. The command iptables -t filter -A INPUT -p tcp --dport 4000:5000 -j ACCEPT will add a new rule to the end of the INPUT chain that will accept the incoming TCP packets that have a destination port between 4000 and 5000. This command will enforce the rule and allow the traffic to the specified ports. This is the correct command to use to accomplish the task. The other options are incorrect because they either use the wrong options (-f instead of -t or -D instead of -A) or do not exist (iptables filter -A INPUT -p tcp --dport 4000:5000 -D ACCEPT or iptables filter -S INPUT -p tcp --dport 4000:5000 -A ACCEPT). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 543.

QUESTION 48
A Linux administrator needs to determine whether a hostname is in the DNS. Which of the following would supply the information that is needed?

A. nslookup
B. rsyn?
C. netstat
D. host

Correct Answer: A
Section:
Explanation:
The commands nslookup or host can be used to determine whether a hostname is in the DNS. The DNS is the domain name system, which is a service that translates domain names into IP addresses and vice versa. The nslookup command is a tool for querying the DNS and obtaining information about a domain name or an IP address. The host command is a similar tool that performs DNS lookups. Both commands can be used to check if a hostname is in the DNS by providing the hostname as an argument and seeing if the command returns a valid IP address or an error message.
For example, the command nslookup www.google.com or host www.google.com will return the IP address of the Google website, while the command nslookup www.nosuchdomain.com or host www.nosuchdomain.com will return an error message indicating that the hostname does not exist.
These commands will supply the information that is needed to determine whether a hostname is in the DNS. These are the correct commands to use for this task. The other options are incorrect because they do not query the DNS or obtain information about a hostname (rsync or netstat). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12:
Managing Network Connections, page 378.

QUESTION 49
A server is experiencing intermittent connection issues. Some connections to the Internet work as intended, but some fail as if there is no connectivity. The systems administrator inspects the server configuration:

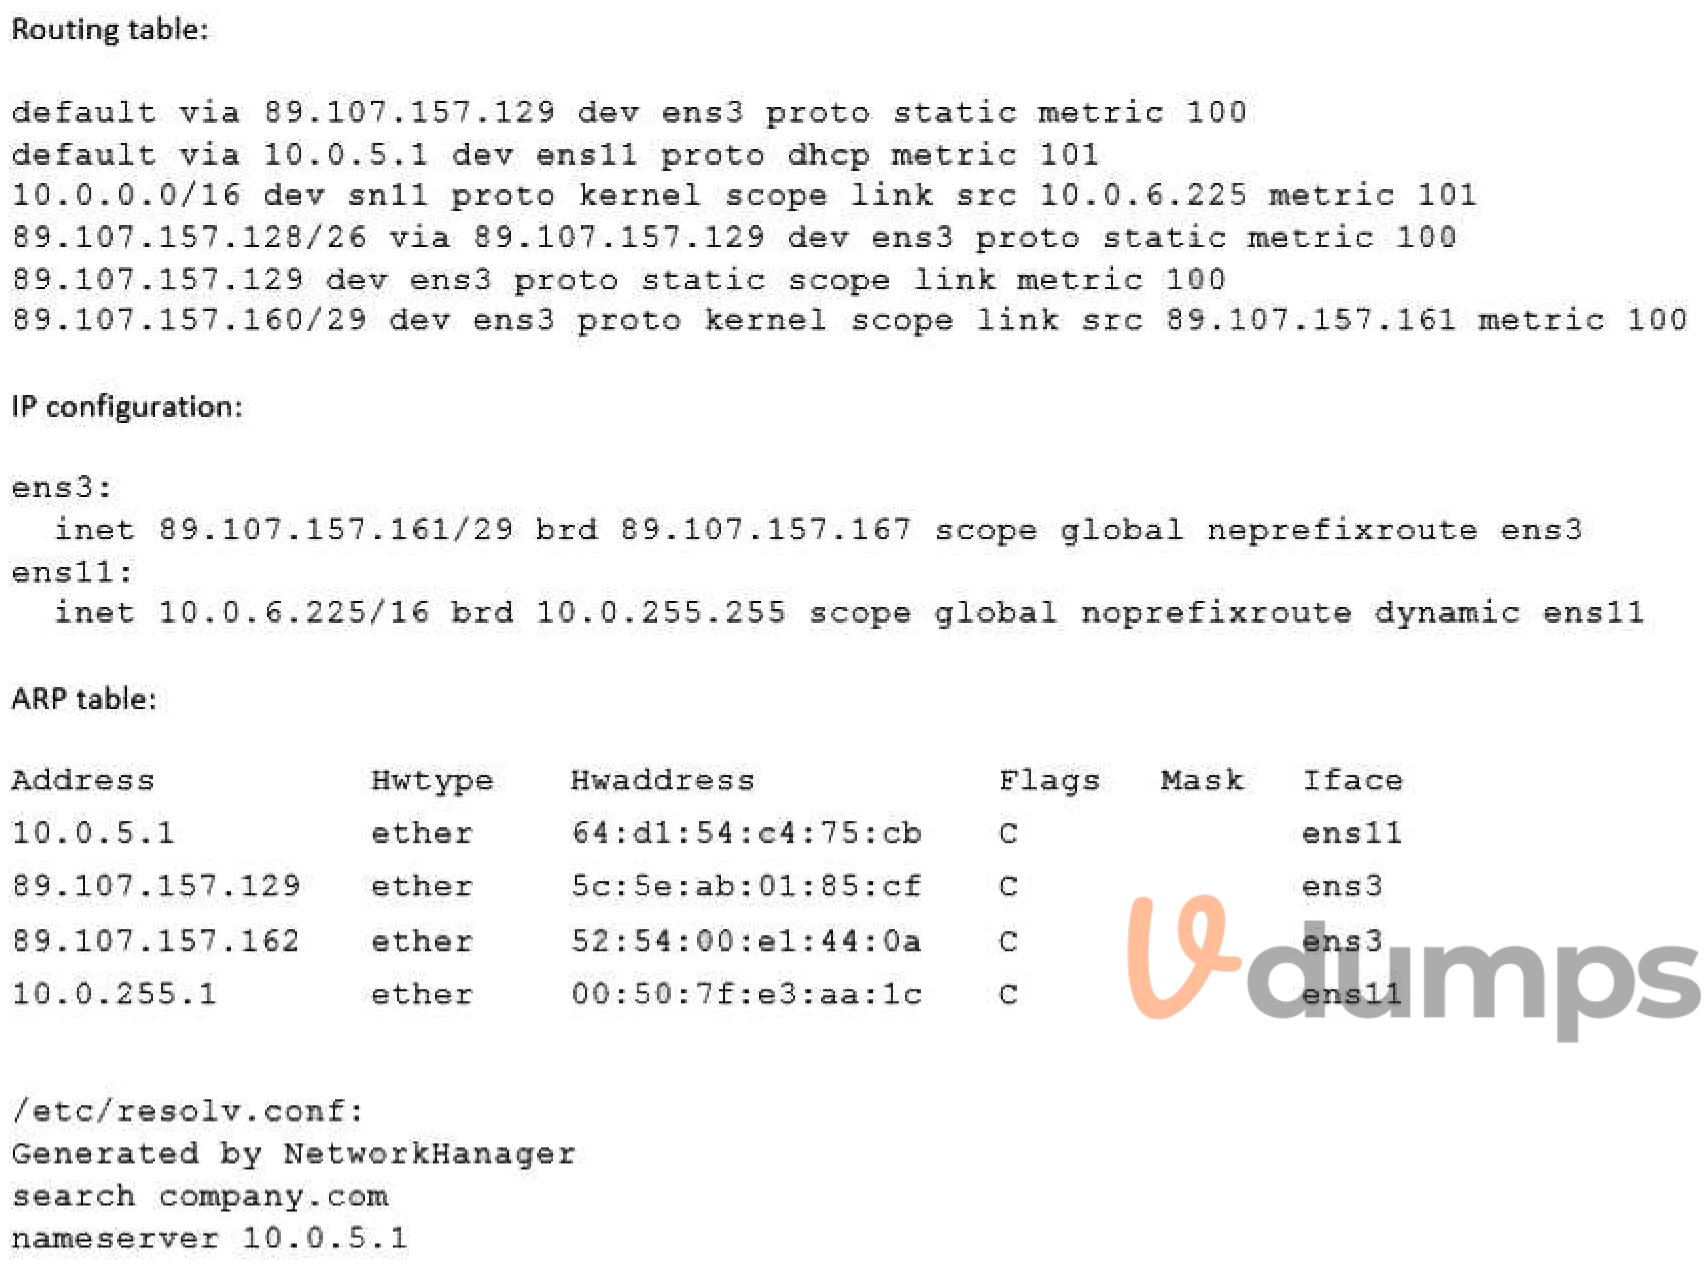

Routing table:

```
default via 89.107.157.129 dev ens3 proto static metric 100
default via 10.0.5.1 dev ens11 proto dhcp metric 101
10.0.0.0/16 dev sn11 proto kernel scope link src 10.0.6.225 metric 101
89.107.157.128/26 via 89.107.157.129 dev ens3 proto static metric 100
89.107.157.129 dev ens3 proto static scope link metric 100
89.107.157.160/29 dev ens3 proto kernel scope link src 89.107.157.161 metric 100
```

IP configuration:

```
ens3:
  inet 89.107.157.161/29 brd 89.107.157.167 scope global neprefixroute ens3
ens11:
  inet 10.0.6.225/16 brd 10.0.255.255 scope global noprefixroute dynamic ens11
```

ARP table:

```
Address          Hwtype   Hwaddress           Flags   Mask   Iface
10.0.5.1         ether    64:d1:54:c4:75:cb   C              ens11
89.107.157.129   ether    5c:5e:ab:01:85:cf   C              ens3
89.107.157.162   ether    52:54:00:e1:44:0a   C              ens3
10.0.255.1       ether    00:50:7f:e3:aa:1c   C              ens11
```

```
/etc/resolv.conf:
Generated by NetworkHanager
search company.com
nameserver 10.0.5.1
```

Which of the following is MOST likely the cause of the issue?

A. An internal-only DNS server is configured.
B. The IP netmask is wrong for ens3.
C. Two default routes are configured.
D. The ARP table contains incorrect entries.

Correct Answer: C
Section:
Explanation:
The most likely cause of the issue is that two default routes are configured on the server. The default route is the route that is used when no other route matches the destination of a packet. The default route is usually the gateway that connects the local network to the Internet. The server configuration shows that there are two default routes in the routing table, one with the gateway 192.168.1.1 and the other with the gateway 10.0.0.1. This can cause a conflict and confusion for the server when deciding which gateway to use for the outgoing packets. Some packets may be sent to the wrong gateway and fail to reach the Internet, while some packets may be sent to the correct gateway and work as intended. This can result in intermittent connection issues and inconsistent behavior. The administrator should remove one of the default routes and keep only the correct one for the network. This can be done by using the ip route del command or by editing the network configuration files. This will resolve the issue and restore the connectivity. The other options are incorrect because they are not supported by the outputs. The DNS server, the IP netmask, and the ARP table are not the causes of the issue. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections,

pages 381-382.

QUESTION 50
A cloud engineer needs to block the IP address 192.168.10.50 from accessing a Linux server. Which of the following commands will achieve this goal?

A. iptables -F INPUT -j 192.168.10.50 -m DROP
B. iptables -A INPUT -s 192.168.10.30 -j DROP
C. iptables -i INPUT --ipv4 192.168.10.50 -z DROP
D. iptables -j INPUT 192.168.10.50 -p DROP

Correct Answer: B
Section:
Explanation:
The correct command to block the IP address 192.168.10.50 from accessing a Linux server is iptables -A INPUT -s 192.168.10.50 -j DROP. This command appends a rule to the INPUT chain that matches the source address 192.168.10.50 and jumps to the DROP target, which discards the packet. The other commands are incorrect because they either have invalid syntax, wrong parameters, or wrong order of arguments. Reference: CompTIA Linux+ Study Guide, Fourth Edition, page 457-458.

QUESTION 51
A Linux systems administrator is configuring a new filesystem that needs the capability to be mounted persistently across reboots. Which of the following commands will accomplish this task? (Choose two.)

A. df -h /data
B. mkfs.ext4 /dev/sdc1
C. fsck /dev/sdc1
D. fdisk -l /dev/sdc1
E. echo "/data /dev/sdc1 ext4 defaults 0 0" >> /etc/fstab
F. echo "/dev/sdc1 /data ext4 defaults 0 0" >> /etc/fstab

Correct Answer: B, F
Section:
Explanation:
"modify the /etc/fstab text file to automatically mount the new partition by opening it in an editor and adding the following line:
/dev/ xxx 1 /data ext4 defaults 1 2 where xxx is the device name of the storage device"
https://learning.oreilly.com/library/view/mastering-linux-system/9781119794455/b01.xhtml
To configure a new filesystem that needs the capability to be mounted persistently across reboots, two commands are needed: mkfs.ext4 /dev/sdc1 and echo "/dev/sdc1 /data ext4 defaults 0 0" >> /etc/fstab. The first command creates an ext4 filesystem on the device /dev/sdc1, which is the partition that will be used for the new filesystem. The second command appends a line to the /etc/fstab file, which is the configuration file that controls persistent mount points of filesystems. The line specifies the device name, the mount point (/data), the filesystem type (ext4), the mount options (defaults), and the dump and pass values (0 0). The other commands are incorrect because they either do not create or configure a filesystem, or they have wrong syntax or arguments.
Reference: CompTIA Linux+ Study Guide, Fourth Edition, page 409-410, 414-415.

QUESTION 52
A Linux administrator is alerted to a storage capacity issue on a server without a specific mount point or directory. Which of the following commands would be MOST helpful for troubleshooting? (Choose two.)

A. parted
B. df
C. mount
D. du
E. fdisk

F. dd
G. ls

Correct Answer: B, D
Section:
Explanation:
To troubleshoot a storage capacity issue on a server without a specific mount point or directory, two commands that would be most helpful are df and du. The df command displays information about disk space usage on all mounted filesystems, including their size, used space, available space, and percentage of usage. The du command displays disk space usage by files and directories in a given path, which can help identify large files or directories that may be taking up too much space. The other commands are incorrect because they either do not show disk space usage, or they are used for other purposes such as partitioning, formatting, checking, mounting, copying, or listing files.
Reference: CompTIA Linux+ Study Guide, Fourth Edition, page 417-419.

QUESTION 53
A systems administrator pressed Ctrl+Z after starting a program using the command line, and the shell prompt was presented. In order to go back to the program, which of the following commands can the administrator use?

A. fg
B. su
C. bg
D. ed

Correct Answer: A
Section:
Explanation:
Ctrl+Z suspended the process, and "fg" will bring it back into the foreground of the shell A Comprehensive and Detailed To go back to a program that was suspended by pressing Ctrl+Z in the command line, the command that can be used is fg. The fg command stands for foreground, and it resumes the job that is next in the queue and brings it to the foreground. Alternatively, if there are more than one suspended jobs, fg can be followed by a job number to resume a specific job. The other commands are incorrect because they either do not resume a suspended job, or they have different functions such as switching user (su), pushing a job to the background (bg), or editing a file (ed). Reference: CompTIA Linux+ Study Guide, Fourth Edition, page 181-182.

QUESTION 54
A systems administrator received a notification that a system is performing slowly. When running the top command, the systems administrator can see the following values:

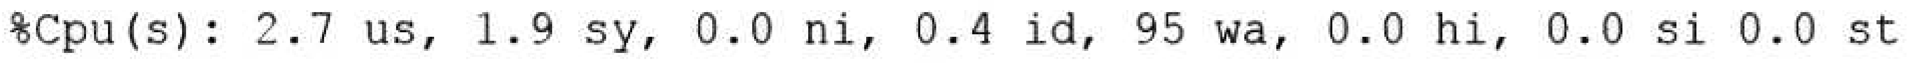

```
%Cpu(s): 2.7 us, 1.9 sy, 0.0 ni, 0.4 id, 95 wa, 0.0 hi, 0.0 si 0.0 st
```

Which of the following commands will the administrator most likely run NEXT?

A. vmstat
B. strace
C. htop
D. lsof

Correct Answer: A
Section:
Explanation:
The command vmstat will most likely be run next by the administrator to troubleshoot the system performance. The vmstat command is a tool for reporting virtual memory statistics on Linux systems. The command shows information about processes, memory, paging, block IO, interrupts, and CPU activity. The command can help the administrator identify the source of the performance issue, such as high CPU usage, low free memory, excessive swapping, or disk IO bottlenecks. The command can also be used with an interval and a count to display the statistics repeatedly over time and observe the changes. The command vmstat will provide useful information for diagnosing the system performance and finding the root cause of the issue. This is the most likely command to run next after the top command. The other options are incorrect because they either do not show the virtual memory statistics (strace or lsof) or do not provide more information than the top command (htop). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, page 425.

QUESTION 55

Which of the following technologies provides load balancing, encryption, and observability in containerized environments?

A. Virtual private network
B. Sidecar pod
C. Overlay network
D. Service mesh

Correct Answer: D
Section:
Explanation:
"A service mesh controls the delivery of service requests in an application. Common features provided by a service mesh include service discovery, load balancing, encryption and failure recovery."
The technology that provides load balancing, encryption, and observability in containerized environments is service mesh. A service mesh is a dedicated infrastructure layer that manages the communication and security between microservices in a distributed system. A service mesh consists of two components: a data plane and a control plane. The data plane is composed of proxies that are deployed alongside the microservices as sidecar pods. The proxies handle the network traffic between the microservices and provide features such as load balancing, encryption, authentication, authorization, routing, and observability. The control plane is responsible for configuring and managing the data plane and providing a unified interface for the administrators and developers. A service mesh can help improve the performance, reliability, and security of containerized applications and simplify the development and deployment process. A service mesh is the technology that provides load balancing, encryption, and observability in containerized environments. This is the correct answer to the question. The other options are incorrect because they either do not provide all the features of a service mesh (virtual private network or overlay network) or are not a technology but a component of a service mesh (sidecar pod).
Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 574.
https://www.techtarget.com/searchitoperations/definition/service-mesh

QUESTION 56

A development team asks an engineer to guarantee the persistency of journal log files across system reboots. Which of the following commands would accomplish this task?

A. grep -i auto /etc/systemd/journald.conf && systemctl restart systemd-journald.service
B. cat /etc/systemd/journald.conf | awk '(print $1,$3)'
C. sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i 'persistent/s/^#//q' /etc/systemd/journald.conf
D. journalctl --list-boots && systemctl restart systemd-journald.service

Correct Answer: C
Section:
Explanation:
The command sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i 'persistent/s/^#//q' /etc/systemd/journald.conf will accomplish the task of guaranteeing the persistency of journal log files across system reboots. The sed command is a tool for editing text files on Linux systems. The -i option modifies the file in place. The s command substitutes one string for another. The g flag replaces all occurrences of the string. The && operator executes the second command only if the first command succeeds. The q command quits after the first match. The /etc/systemd/journald.conf file is a configuration file for the systemd-journald service, which is responsible for collecting and storing log messages. The command sed -i 's/auto/persistent/g' /etc/systemd/journald.conf will replace the word auto with the word persistent in the file. This will change the value of the Storage option, which controls where the journal log files are stored. The value auto means that the journal log files are stored in the volatile memory and are lost after reboot, while the value persistent means that the journal log files are stored in the persistent storage and are preserved across reboots. The command sed -i 'persistent/s/^#//q' /etc/systemd/journald.conf will remove the # character at the beginning of the line that contains the word persistent. This will uncomment the Storage option and enable it. The command sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i 'persistent/s/^#//q' /etc/systemd/journald.conf will guarantee the persistency of journal log files across system reboots by changing and enabling the Storage option to persistent. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not change the value of the Storage option (grep -i auto /etc/systemd/journald.conf && systemctl restart systemd-journald.service or cat /etc/systemd/journald.conf | awk '(print $1,$3)') or do not enable the Storage option (journalctl --list-boots && systemctl restart systemdjournald. service). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16:
Managing Logging and Monitoring, page 489.

QUESTION 57

A systems administrator is receiving tickets from users who cannot reach the application app that should be listening on port 9443/tcp on a Linux server.
To troubleshoot the issue, the systems administrator runs netstat and receives the following output:

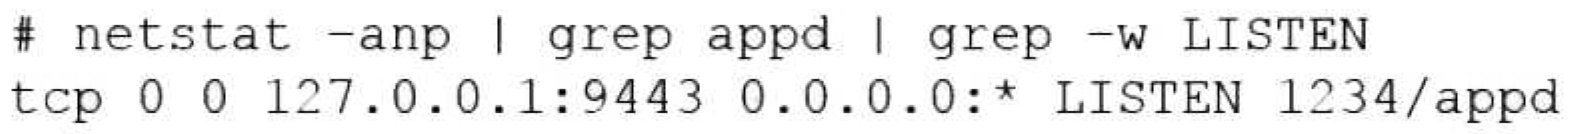

```
# netstat -anp | grep appd | grep -w LISTEN
tcp 0 0 127.0.0.1:9443 0.0.0.0:* LISTEN 1234/appd
```

Based on the information above, which of the following is causing the issue?

A. The IP address 0.0.0.0 is not valid.
B. The application is listening on the loopback interface.
C. The application is listening on port 1234.
D. The application is not running.

Correct Answer: B
Section:
Explanation:
The server is in a "Listen" state on port 9943 using its loopback address. The "1234" is a process-id
The cause of the issue is that the application is listening on the loopback interface. The loopback interface is a virtual network interface that is used for internal communication within the system.
The loopback interface has the IP address 127.0.0.1, which is also known as localhost. The netstat output shows that the application is listening on port 9443 using the IP address 127.0.0.1. This means that the application can only accept connections from the same system, not from other systems on the network. This can prevent the users from reaching the application and cause the issue. The administrator should configure the application to listen on the IP address 0.0.0.0, which means all available interfaces, or on the specific IP address of the system that is reachable from the network.
This will allow the application to accept connections from other systems and resolve the issue. The cause of the issue is that the application is listening on the loopback interface. This is the correct answer to the question. The other options are incorrect because they are not supported by the outputs. The IP address 0.0.0.0 is valid and means all interfaces, the application is not listening on port 1234, and the application is running as shown by the process ID 1234. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 383.

QUESTION 58
A systems administrator is troubleshooting a connectivity issue pertaining to access to a system named db.example.com. The system IP address should be 192.168.20.88. The administrator issues the dig command and receives the following output:

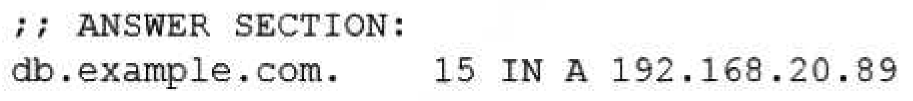

```
;; ANSWER SECTION:
db.example.com.    15 IN A 192.168.20.89
```

The administrator runs grep db.example.com /etc/hosts and receives the following output:

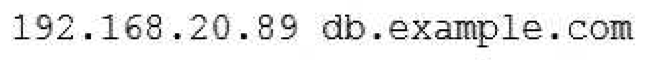

```
192.168.20.89 db.example.com
```

Given this scenario, which of the following should the administrator do to address this issue?

A. Modify the /etc/hosts file and change the db.example.com entry to 192.168.20.89.
B. Modify the /etc/network file and change the db.example.com entry to 192.168.20.88.
C. Modify the /etc/network file and change the db.example.com entry to 192.168.20.89.
D. Modify the /etc/hosts file and change the db.example.com entry to 192.168.20.88.

Correct Answer: D
Section:
Explanation:
The administrator should modify the /etc/hosts file and change the db.example.com entry to 192.168.20.88 to address the issue. The /etc/hosts file is a file that maps hostnames to IP addresses on Linux systems. The file can be used to override the DNS resolution and provide a local lookup for hostnames. The dig output shows that the DNS returns the IP address 192.168.20.88 for the hostname db.example.com, which is the correct IP address of the system. The grep output shows that the /etc/hosts file contains an entry for db.example.com with the IP address 192.168.20.89, which is the wrong IP address of the system. This can cause a conflict and prevent the system from being accessed by the hostname. The administrator should modify the /etc/hosts file and change the db.example.com entry to 192.168.20.88, which is the correct IP address of the system. This will align the

/etc/hosts file with the DNS and allow the system to be accessed by the hostname. The administrator should modify the /etc/hosts file and change the db.example.com entry to 192.168.20.88 to address the issue. This is the correct answer to the question. The other options are incorrect because they either do not modify the /etc/hosts file (modify the /etc/network file and change the db.example.com entry to 192.168.20.88 or modify the /etc/network file and change the db.example.com entry to 192.168.20.89) or do not change the IP address to the correct one (modify the /etc/hosts file and change the db.example.com entry to 192.168.20.89). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 378.

QUESTION 59

Users have been unable to reach www.comptia.org from a Linux server. A systems administrator is troubleshooting the issue and does the following:

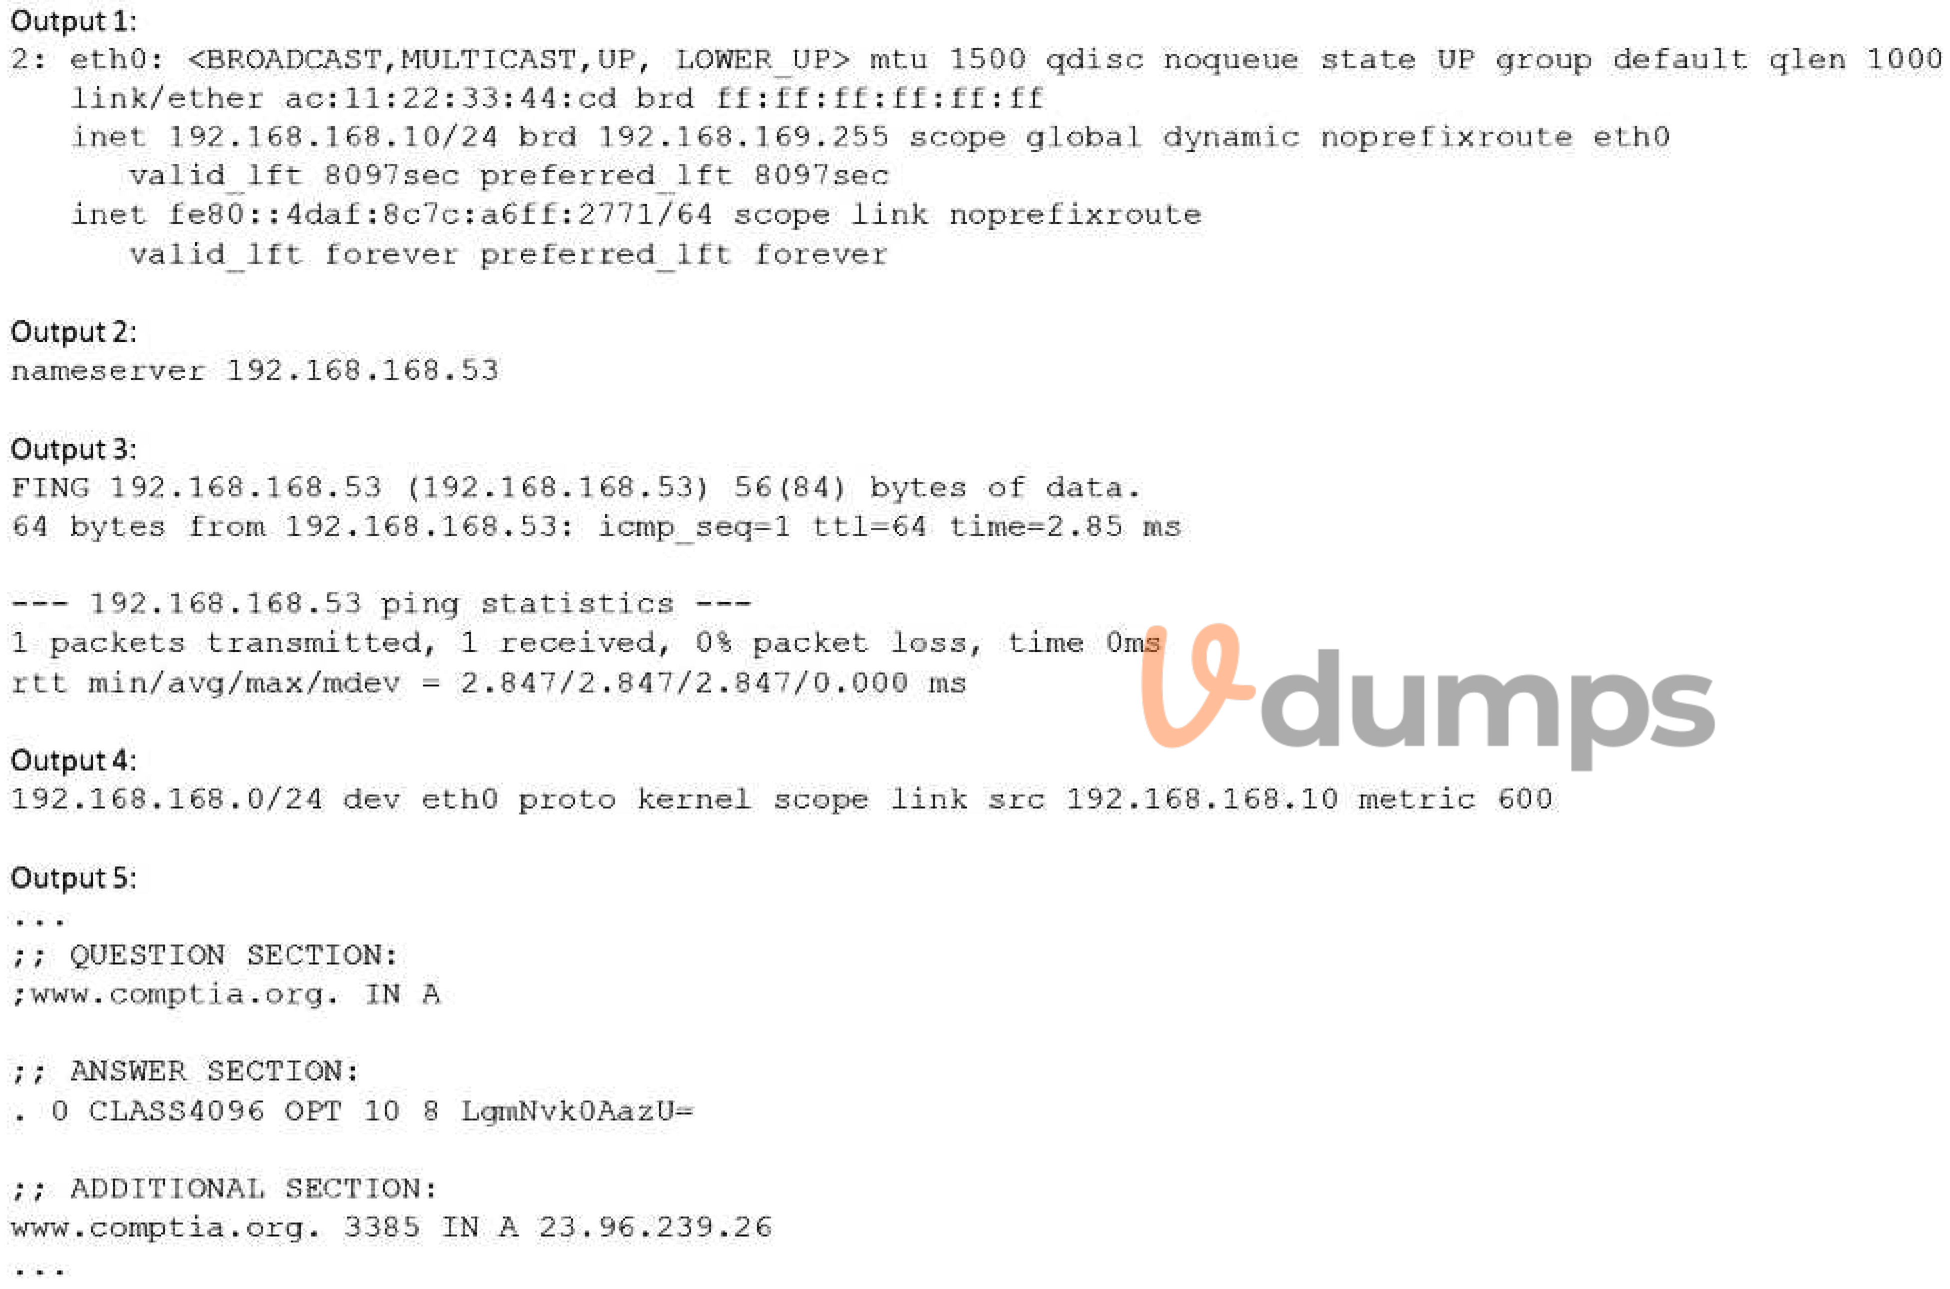

Output 1:

```
2: eth0: <BROADCAST,MULTICAST,UP, LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
   link/ether ac:11:22:33:44:cd brd ff:ff:ff:ff:ff:ff
   inet 192.168.168.10/24 brd 192.168.169.255 scope global dynamic noprefixroute eth0
      valid_lft 8097sec preferred_lft 8097sec
   inet fe80::4daf:8c7c:a6ff:2771/64 scope link noprefixroute
      valid_lft forever preferred_lft forever
```

Output 2:

```
nameserver 192.168.168.53
```

Output 3:

```
FING 192.168.168.53 (192.168.168.53) 56(84) bytes of data.
64 bytes from 192.168.168.53: icmp_seq=1 ttl=64 time=2.85 ms

--- 192.168.168.53 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 2.847/2.847/2.847/0.000 ms
```

Output 4:

```
192.168.168.0/24 dev eth0 proto kernel scope link src 192.168.168.10 metric 600
```

Output 5:

```
...
;; QUESTION SECTION:
;www.comptia.org. IN A

;; ANSWER SECTION:
. 0 CLASS4096 OPT 10 8 LgmNvk0AazU=

;; ADDITIONAL SECTION:
www.comptia.org. 3385 IN A 23.96.239.26
...
```

Based on the information above, which of the following is causing the issue?

A. The name www.comptia.org does not point to a valid IP address.
B. The server 192.168.168.53 is unreachable.
C. No default route is set on the server.
D. The network interface eth0 is disconnected.

Correct Answer: B

Section:
Explanation:
The issue is caused by the server 192.168.168.53 being unreachable. This server is the DNS server configured in the /etc/resolv.conf file, which is used to resolve domain names to IP addresses. The ping command shows that the server cannot be reached, and the nslookup command shows that the name www.comptia.org cannot be resolved using this server. The other options are incorrect because:
The name www.comptia.org does point to a valid IP address, as shown by the nslookup command using another DNS server (8.8.8.8).
The default route is set on the server, as shown by the ip route command, which shows a default gateway of 192.168.168.1.
The network interface eth0 is connected, as shown by the ip link command, which shows a state of UP for eth0. Reference: CompTIA Linux+ Study Guide, Fourth Edition, page 457-458, 461-462.

QUESTION 60
A systems technician is working on deploying several microservices to various RPM-based systems, some of which could run up to two hours. Which of the following commands will allow the technician to execute those services and continue deploying other microservices within the same terminal section?

A. gedit & disown
B. kill 9 %1
C. fg %1
D. bg %1 job name

Correct Answer: D
Section:
Explanation:
The command that will allow the technician to execute the services and continue deploying other microservices within the same terminal session is bg %1 job name. This command will send the job with ID 1 and name job name to the background, where it will run without occupying the terminal.
The other options are incorrect because:
gedit & disown will launch a graphical text editor in the background and detach it from the terminal, but it will not execute any service.
kill 9 %1 will terminate the job with ID 1 using a SIGKILL signal, which cannot be ignored or handled by the process.
fg %1 will bring the job with ID 1 to the foreground, where it will occupy the terminal until it finishes or is stopped. Reference: CompTIA Linux+ Study Guide, Fourth Edition, page 181-182.

QUESTION 61
A Linux administrator was notified that a virtual server has an I/O bottleneck. The Linux administrator analyzes the following output:

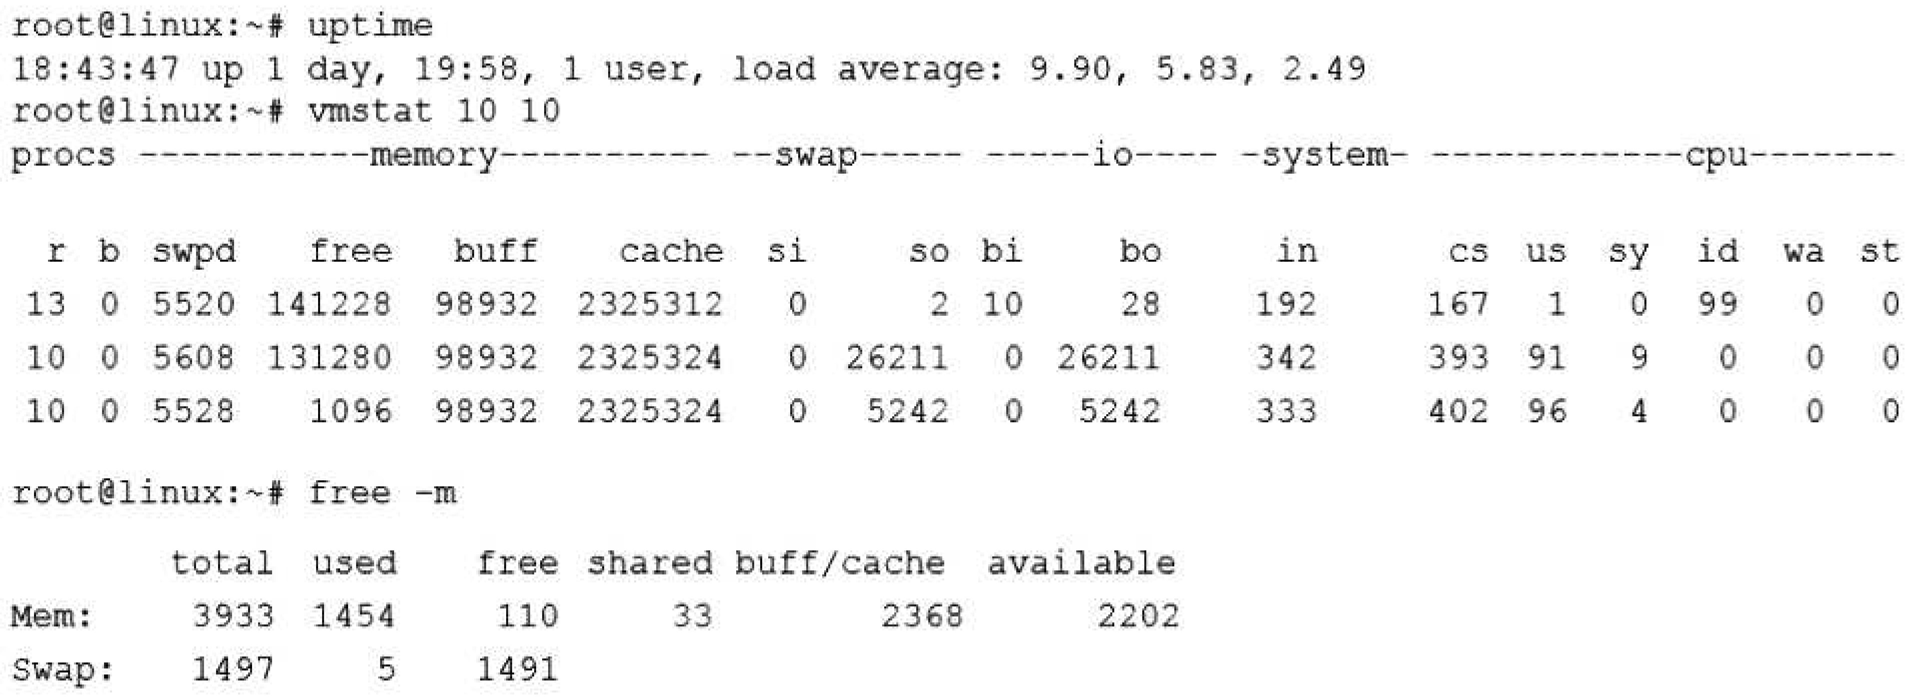

```
root@linux:~# uptime
18:43:47 up 1 day, 19:58, 1 user, load average: 9.90, 5.83, 2.49
root@linux:~# vmstat 10 10
procs -----------memory---------- --swap----- -----io---- -system- ------------cpu-------

  r  b  swpd   free   buff   cache  si     so  bi    bo    in     cs  us  sy  id  wa  st
 13  0  5520 141228  98932 2325312   0      2  10    28   192    167   1   0  99   0   0
 10  0  5608 131280  98932 2325324   0  26211   0 26211   342    393  91   9   0   0   0
 10  0  5528   1096  98932 2325324   0   5242   0  5242   333    402  96   4   0   0   0

root@linux:~# free -m

         total  used   free  shared buff/cache  available
Mem:      3933  1454    110      33       2368       2202
Swap:     1497     5   1491
```

Given there is a single CPU in the sever, which of the following is causing the slowness?

A. The system is running out of swap space.
B. The CPU is overloaded.
C. The memory is exhausted.
D. The processes are paging.

Correct Answer: B
Section:
Explanation:
The slowness is caused by the CPU being overloaded. The iostat command shows that the CPU utilization is 100%, which means that there are more processes competing for CPU time than the CPU can handle. The other options are incorrect because:
The system is not running out of swap space, as shown by the iostat command, which shows that there is no swap activity (si and so columns are zero).
The memory is not exhausted, as shown by the free -m command, which shows that there is still available memory (avail column) and free buffer/cache memory (buff/cache column).
The processes are not paging, as shown by the vmstat command, which shows that there are no major page faults (majflt column) and no swap activity (si and so columns). Reference: CompTIA Linux+ Study Guide, Fourth Edition, page 417-419, 424-425.

QUESTION 62
Employees in the finance department are having trouble accessing the file /opt/work/file. All IT employees can read and write the file. Systems administrator reviews the following output:

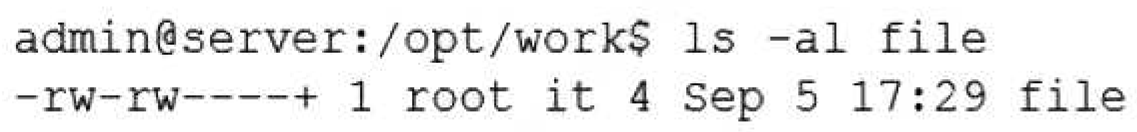

```
admin@server:/opt/work$ ls -al file
-rw-rw----+ 1 root it 4 Sep 5 17:29 file
```

Which of the following commands would permanently fix the access issue while limiting access to IT and finance department employees?

A. chattr +i file
B. chown it:finance file
C. chmod 666 file
D. setfacl -m g:finance:rw file

Correct Answer: D
Section:
Explanation:
The command setfacl -m g:finance:rw file will permanently fix the access issue while limiting access to IT and finance department employees. The setfacl command is a tool for modifying the access control lists (ACLs) of files and directories on Linux systems. The ACLs are a mechanism that allows more fine-grained control over the permissions of files and directories than the traditional ownergroup-others model. The -m option specifies the modification to the ACL. The g:finance:rw means that the group named finance will have read and write permissions on the file. The file is the name of the file to modify, in this case /opt/work/file. The command setfacl -m g:finance:rw file will add an entry to the ACL of the file that will grant read and write access to the finance group. This will fix the access issue and allow the finance employees to access the file. The command will also preserve the existing permissions of the file, which means that the IT employees will still have read and write access to the file. This will limit the access to IT and finance department employees and prevent unauthorized access from other users. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not fix the access issue (chattr +i file or chown it:finance file) or do not limit the access to IT and finance department employees (chmod 666 file). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing File Permissions and Ownership, page 352.

QUESTION 63
A Linux engineer needs to create a custom script, cleanup.sh, to run at boot as part of the system services. Which of the following processes would accomplish this task?

A.
Create a unit file in the /etc/default/ directory.
systemctl enable cleanup
systemctl is-enabled cleanup

B.

Create a unit file in the /etc/ske1/ directory.
systemctl enable cleanup
systemctl is-enabled cleanup

C.

Create a unit file in the /etc/systemd/system/ directory.
systemctl enable cleanup
systemctl is-enabled cleanup

D.

Create a unit file in the /etc/sysctl.d/ directory.
systemctl enable cleanup
systemctl is-enabled cleanup

Correct Answer: C
Section:
Explanation:
The process that will accomplish the task of creating a custom script to run at boot as part of the system services is:
Create a unit file in the /etc/systemd/system/ directory. A unit file is a configuration file that defines the properties and behavior of a systemd service. The systemd is a system and service manager that controls the startup and operation of Linux systems. The /etc/systemd/system/ directory is the location where the administrator can create and store custom unit files. The unit file should have a name that matches the name of the script, such as cleanup.service, and should contain the following sections and options:
[Unit]: This section provides the general information about the service, such as the description, dependencies, and conditions. The administrator should specify the following options in this section:
Description: A brief description of the service, such as "Custom cleanup script".
After: The name of another unit that this service should start after, such as "network.target".
ConditionPathExists: The path of the file or directory that must exist for the service to start, such as "/opt/scripts/cleanup.sh".
[Service]: This section defines how the service should be started and stopped, and what commands should be executed. The administrator should specify the following options in this section:
Type: The type of the service, such as "oneshot", which means that the service will run once and then exit.
ExecStart: The command that will start the service, such as "/bin/bash /opt/scripts/cleanup.sh".
RemainAfterExit: A boolean value that indicates whether the service should remain active after the command exits, such as "yes".
[Install]: This section defines how the service should be enabled and under what circumstances it should be started. The administrator should specify the following option in this section:
WantedBy: The name of another unit that wants this service to be started, such as "multiuser. target", which means that the service will be started when the system reaches the multi-user mode.
Run the command systemctl enable cleanup. This command will enable the service and create the necessary symbolic links to start the service at boot.
Run the command systemctl is-enabled cleanup. This command will check the status of the service and confirm that it is enabled.
This process will create a custom script, cleanup.sh, to run at boot as part of the system services. This is the correct process to use to accomplish the task. The other options are incorrect because they either use the wrong directory for the unit file (/etc/default/, /etc/skel/, or /etc/sysctl.d/) or do not create a unit file at all. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15:
Managing System Services, pages 457-459.

QUESTION 64
A Linux system is failing to boot. The following error is displayed in the serial console:
[[1;33mDEPEND[Om] Dependency failed for /data.
[[1;33mDEPEND[Om] Dependency failed for Local File Systems
... Welcome to emergency mode! After logging in, type "journalctl -xb" to viewsystem logs, "systemctl reboot" to reboot, "systemctl default" to try again to boot into default mode.
Give root password for maintenance (or type Control-D to continue}
Which of the following files will need to be modified for this server to be able to boot again?

A. /etc/mtab
B. /dev/sda
C. /etc/fstab
D. /ete/grub.conf

Correct Answer: C
Section:
Explanation:
The file that will need to be modified for the server to be able to boot again is /etc/fstab. The /etc/fstab file is a file that contains the information about the file systems that are mounted at boot time on Linux systems. The file specifies the device name, mount point, file system type, mount options, dump frequency, and pass number for each file system. The error message indicates that the dependency failed for /data, which is a mount point for a file system. This means that the system could not mount the /data file system at boot time, which caused the system to enter the emergency mode. The emergency mode is a mode that allows the administrator to log in as the root user and perform basic tasks such as repairing the system. The administrator should modify the /etc/fstab file and check the entry for the /data file system. The administrator should look for any errors or inconsistencies in the device name, file system type, or mount options, and correct them. The administrator should also verify that the device and the file system are intact and functional by using commands such as blkid, fdisk, fsck, or mount. The administrator should then reboot the system and see if the issue is resolved. The file that will need to be modified for the server to be able to boot again is /etc/fstab. This is the correct answer to the question. The other options are incorrect because they are not related to the file systems that are mounted at boot time (/etc/mtab, /dev/sda, or /etc/grub.conf). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10:
Managing Storage, page 321.

QUESTION 65
A systems administrator frequently connects to a remote host via SSH and a non-standard port. The systems administrator would like to avoid passing the port parameter on the command line every time. Which of the following files can be used to set a different port value for that host?

A. /etc/ssh/sshd_config
B. /etc/ssh/moduli
C. ~/.ssh/config
D. ~/.ssh/authorized_keys

Correct Answer: C
Section:
Explanation:
The ~/.ssh/config file can be used to set various options for SSH connections, including the port number, for specific hosts or groups of hosts. This file is located in the user's home directory and affects only the current user. The /etc/ssh/sshd_config file is used to configure the SSH server daemon, not the client. The /etc/ssh/moduli file contains parameters for Diffie-Hellman key exchange, not port settings. The ~/.ssh/authorized_keys file contains public keys for authentication, not port settings. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Secure Shell (SSH), page 414.

QUESTION 66
A Linux administrator modified the SSH configuration file. Which of the following commands should be used to apply the configuration changes?

A. systemctl stop sshd
B. systemctl mask sshd
C. systemctl reload sshd
D. systemctl start sshd

Correct Answer: C
Section:
Explanation:
The systemctl reload sshd command can be used to apply the configuration changes of the SSH server daemon without restarting it. This is useful to avoid interrupting existing connections. The systemctl stop sshd command would stop the SSH server daemon, not apply the changes. The systemctl mask sshd command would prevent the SSH server daemon from being started, not apply the changes. The systemctl start sshd command would start the SSH server daemon if it is not running, but it would not apply the changes if it is already running. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Secure Shell (SSH), page 415.

QUESTION 67
A cloud engineer needs to check the link status of a network interface named eth1 in a Linux server.
Which of the following commands can help to achieve the goal?

A. ifconfig hw eth1
B. netstat -r eth1
C. ss -ti eth1
D. ip link show eth1

Correct Answer: D
Section:
Explanation:
The ip link show eth1 command can be used to check the link status of a network interface named eth1 in a Linux server. It will display information such as the MAC address, MTU, state, and flags of the interface. The ifconfig hw eth1 command is invalid, as hw is not a valid option for ifconfig. The netstat -r eth1 command would display the routing table for eth1, not the link status. The ss -ti eth1 command would display TCP information for sockets associated with eth1, not the link status. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Networking Fundamentals, page 436.

QUESTION 68
A systems administrator is tasked with setting up key-based SSH authentication. In which of the following locations should the administrator place the public keys for the server?

A. ~/.sshd/authkeys
B. ~/.ssh/keys
C. ~/.ssh/authorized_keys
D. ~/.ssh/keyauth

Correct Answer: C
Section:
Explanation:
The administrator should place the public keys for the server in the ~/.ssh/authorized_keys file. The SSH (Secure Shell) protocol is a method for establishing secure and encrypted connections between remote systems. The SSH protocol supports two types of authentication: password-based and keybased.
Password-based authentication requires the user to enter the password of the remote system every time they connect. Key-based authentication requires the user to generate a pair of cryptographic keys: a public key and a private key. The public key is stored on the remote system, while the private key is kept on the local system. The public key and the private key are mathematically related, but not identical. The SSH protocol uses the keys to verify the identity of the user and establish a secure connection without requiring a password. The ~/.ssh/authorized_keys file is a file that contains the public keys of the users who are allowed to connect to the remote system using key-based authentication. The administrator should place the public keys for the server in this file, one per line, and set the appropriate permissions for the file. The administrator should also configure the SSH server to enable key-based authentication by editing the /etc/ssh/sshd_config file and setting the option PasswordAuthentication to no. The administrator should place the public keys for the server in the ~/.ssh/authorized_keys file. This is the correct answer to the question. The other options are incorrect because they are not the standard locations for the public keys for the server (~/.sshd/authkeys, ~/.ssh/keys, or ~/.ssh/keyauth). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Implementing Basic Security, page 513.

QUESTION 69
A systems administrator has been unable to terminate a process. Which of the following should the administrator use to forcibly stop the process?

A. kill -1
B. kill -3
C. kill -15
D. kill -HUP
E. kill -TERM

Correct Answer: E
Section:
Explanation:
The administrator should use the command kill -TERM to forcibly stop the process. The kill command is a tool for sending signals to processes on Linux systems. Signals are messages that inform the processes about certain events and actions. The processes can react to the signals by performing predefined or user-defined actions, such as terminating, suspending, resuming, or ignoring. The -TERM option specifies the signal name or number that the kill command should send. The TERM signal, which stands for terminate, is the default signal that the kill command sends if no option is specified. The TERM signal requests the process to terminate gracefully, by

closing any open files, releasing any resources, and performing any cleanup tasks. However, if the process does not respond to the TERM signal, the kill command can send a stronger signal, such as the KILL signal, which forces the process to terminate immediately, without any cleanup. The administrator should use the command kill -TERM to forcibly stop the process. This is the correct answer to the question. The other options are incorrect because they either do not terminate the process (kill -1 or kill -3) or do not terminate the process forcibly (kill -15 or kill -HUP). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes, page 431.

QUESTION 70

A systems administrator is compiling a report containing information about processes that are listening on the network ports of a Linux server. Which of the following commands will allow the administrator to obtain the needed information?

A. ss -pint

B. tcpdump -nL

C. netstat -pn

D. lsof -lt

Correct Answer: A
Section:
Explanation:

The command ss -pint will allow the administrator to obtain the needed information about processes that are listening on the network ports of a Linux server. The ss command is a tool for displaying socket statistics on Linux systems. Sockets are endpoints of network communication that allow processes to exchange data over the network. The ss command can show various information about the sockets, such as the state, address, port, protocol, and process. The -pint option specifies the filters and flags that the ss command should apply. The -p option shows the process name and ID that owns the socket. The -i option shows the internal information about the socket, such as the send and receive queue, the congestion window, and the retransmission timeout. The -n option shows the numerical address and port, instead of resolving the hostnames and service names. The -t option shows only the TCP sockets, which are the most common type of sockets used for network communication. The command ss -pint will display the socket statistics for the TCP sockets, along with the process name and ID, the numerical address and port, and the internal information. This will allow the administrator to obtain the needed information about processes that are listening on the network ports of a Linux server. This is the correct command to use to obtain the needed information. The other options are incorrect because they either do not show the socket statistics (tcpdump -nL or lsof -lt) or do not show the process name and ID (netstat -pn). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 389.

QUESTION 71

User1 is a member of the accounting group. Members of this group need to be able to execute but not make changes to a script maintained by User2. The script should not be accessible to other users or groups. Which of the following will give proper access to the script?

A.

```
chown user2:accounting script.sh
chmod 750 script.sh
```

B.

```
chown user1:accounting script.sh
chmod 777 script.sh
```

C.

```
chown accounting:user1 script.sh
chmod 057 script.sh
```

D.

```
chown user2:accounting script.sh
chmod u+x script.sh
```

Correct Answer: A
Section:
Explanation:

The commands that will give proper access to the script are:

chown user2:accounting script.sh: This command will change the ownership of the script to user2 as the owner and accounting as the group. The chown command is a tool for changing the owner and group of files and directories on Linux systems. The user2:accounting is the user and group name that the command should assign to the script. The script.sh is the name of the script that the command should modify. The command chown

user2:accounting script.sh will ensure that user2 is the owner of the script and accounting is the group of the script, which will allow user2 to maintain the script and the accounting group to access the script.
chmod 750 script.sh: This command will change the permissions of the script to 750, which means read, write, and execute for the owner; read and execute for the group; and no access for others.
The chmod command is a tool for changing the permissions of files and directories on Linux systems.
The permissions are represented by three digits in octal notation, where each digit corresponds to the owner, group, and others. Each digit can have a value from 0 to 7, where each value represents a combination of read, write, and execute permissions. The 750 is the permission value that the command should assign to the script. The script.sh is the name of the script that the command should modify. The command chmod 750 script.sh will ensure that only the owner and the group can execute the script, but not make changes to it, and that the script is not accessible to other users or groups.
The commands that will give proper access to the script are chown user2:accounting script.sh and chmod 750 script.sh. This is the correct answer to the question. The other options are incorrect because they either do not give proper access to the script (chown user1:accounting script.sh or chown accounting:user1 script.sh) or do not change the permissions of the script (chmod 777 script.sh or chmod u+x script.sh). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing File Permissions and Ownership, pages 346-348.

QUESTION 72
A systems administrator needs to verify whether the built container has the app.go file in its root directory. Which of the following can the administrator use to verify the root directory has this file?

A. docker image inspect
B. docker container inspect
C. docker exec <container_name> ls
D. docker ps <container_name>

Correct Answer: C
Section:
Explanation:
The docker exec <container_name> ls command can be used to verify whether the built container has the app.go file in its root directory. This command will run the ls command inside the specified container and list the files and directories in its root directory. If the app.go file is present, it will be displayed in the output. The docker image inspect command will display information about an image, not a container, and it will not list the files inside the image. The docker container inspect command will display information about a container, not its files. The docker ps <container_name> command is invalid, as ps does not accept a container name as an argument.
Reference: CompTIA
Linux+ (XK0-005) Certification Study Guide, Chapter 16: Virtualization and Cloud Technologies, page 499.

QUESTION 73
Joe, a user, is unable to log in to the Linux system. Given the following output:

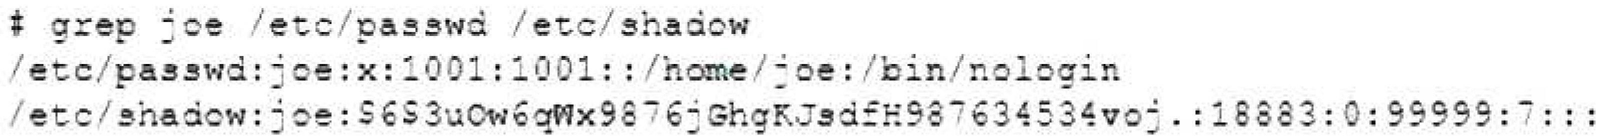

```
# grep joe /etc/passwd /etc/shadow
/etc/passwd:joe:x:1001:1001::/home/joe:/bin/nologin
/etc/shadow:joe:$6$3uOw6qWx9876jGhgKJsdfH987634534voj.:18883:0:99999:7:::
```

Which of the following commands would resolve the issue?

A. usermod -s /bin/bash joe
B. pam_tally2 -u joe -r
C. passwd -u joe
D. chage -E 90 joe

Correct Answer: B
Section:
Explanation:
The command pam_tally2 -u joe -r will resolve the issue of Joe being unable to log in to the Linux system. The pam_tally2 command is a tool for managing the login counter for the PAM (Pluggable Authentication Modules) system. PAM is a framework for managing authentication and authorization on Linux systems. PAM allows the administrator to define the rules and policies for accessing various system resources and services, such as login, sudo, ssh, or cron. PAM also supports different types of authentication methods, such as passwords, tokens, biometrics, or smart cards. PAM can be used to implement login restrictions, such as limiting the number of failed login attempts, locking the account after a certain number of failures, or enforcing a minimum or maximum time between login attempts. The pam_tally2 command can display, reset, or unlock the login counter for the users or hosts. The -u joe option specifies the user name that the command should apply to. The -r option resets the login counter for the user. The command pam_tally2 -u joe -r will reset the login counter for Joe, which will unlock his account and allow him to log in to the Linux system. This will resolve the issue of Joe being unable to log in to the Linux system. This is the correct command to use to resolve the issue. The other options are

incorrect because they either do not unlock the account (usermod -s /bin/bash joe or passwd -u joe) or do not affect the login counter (chage -E 90 joe). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Implementing Basic Security, page 517.

QUESTION 74

Users have been unable to save documents to /home/tmp/temp and have been receiving the following error:
Path not found
A junior technician checks the locations and sees that /home/tmp/tempa was accidentally created instead of /home/tmp/temp. Which of the following commands should the technician use to fix this issue?

A. cp /home/tmp/tempa /home/tmp/temp
B. mv /home/tmp/tempa /home/tmp/temp
C. cd /temp/tmp/tempa
D. ls /home/tmp/tempa

Correct Answer: B
Section:
Explanation:
The mv /home/tmp/tempa /home/tmp/temp command will fix the issue of the misnamed directory.
This command will rename the directory /home/tmp/tempa to /home/tmp/temp, which is the expected path for users to save their documents. The cp /home/tmp/tempa /home/tmp/temp command will not fix the issue, as it will copy the contents of /home/tmp/tempa to a new file named /home/tmp/temp, not a directory. The cd /temp/tmp/tempa command will not fix the issue, as it will change the current working directory to /temp/tmp/tempa, which does not exist. The ls /home/tmp/tempa command will not fix the issue, as it will list the contents of /home/tmp/tempa, not rename it. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12:
Managing Files and Directories, page 413.

QUESTION 75

A database administrator requested the installation of a custom database on one of the servers.
Which of the following should the Linux administrator configure so the requested packages can be installed?

A. /etc/yum.conf
B. /etc/ssh/sshd.conf
C. /etc/yum.repos.d/db.repo
D. /etc/resolv.conf

Correct Answer: C
Section:
Explanation:
The Linux administrator should configure /etc/yum.repos.d/db.repo so that the requested packages can be installed. This file defines a custom repository for yum, which is a package manager for RPMbased systems. The file should contain information such as the name, baseurl, gpgcheck, and enabled options for the repository. By creating this file and enabling the repository, the administrator can use yum to install packages from the custom repository. The /etc/yum.conf file is the main configuration file for yum, but it does not define repositories. The /etc/ssh/sshd.conf file is the configuration file for sshd, which is a daemon that provides secure shell access to remote systems.
The /etc/resolv.conf file is the configuration file for DNS resolution, which maps domain names to IP addresses. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Packages and Software, page 559.

QUESTION 76

At what point is the Internal Certificate Authority (ICA) created?

A. During the primary Security Management Server installation process.
B. Upon creation of a certificate.
C. When an administrator decides to create one.

D. When an administrator initially logs into SmartConsole.

Correct Answer: A
Section:
Explanation:
The Internal Certificate Authority (ICA) is created during the primary Security Management Server installation process. The ICA is a component of Check Point's Public Key Infrastructure (PKI) that issues and manages certificates for Security Gateways and administrators. The ICA is automatically installed and initialized when the primary Security Management Server is installed. The ICA is not created upon creation of a certificate, when an administrator decides to create one, or when an administrator initially logs into SmartConsole. Reference: Check Point Certified Security Administrator (CCSA) R80.x Study Guide, Chapter 3: Check Point Security Management Architecture, page 32.

QUESTION 77
Rugged appliances are small appliances with ruggedized hardware and like Quantum Spark appliance they use which operating system?

A. Centos Linux
B. Gaia embedded
C. Gaia
D. Red Hat Enterprise Linux version 5

Correct Answer: B
Section:
Explanation:
Rugged appliances are small appliances with ruggedized hardware that use Gaia embedded as their operating system. Gaia embedded is a version of Gaia that is optimized for embedded devices such as Rugged appliances and Quantum Spark appliances. Gaia embedded supports features such as VPN, firewall, identity awareness, application control, URL filtering, and anti-bot. Gaia embedded does not use Centos Linux, Gaia, or Red Hat Enterprise Linux version 5 as their operating system. Reference: Check Point Rugged Appliance Datasheet, page 1.

QUESTION 78
Using AD Query, the security gateway connections to the Active Directory Domain Controllers using what protocol?

A. Windows Management Instrumentation (WMI)
B. Hypertext Transfer Protocol Secure (HTTPS)
C. Lightweight Directory Access Protocol (LDAP)
D. Remote Desktop Protocol (RDP)

Correct Answer: C
Section:
Explanation:
Using AD Query, the security gateway connects to the Active Directory Domain Controllers using Lightweight Directory Access Protocol (LDAP). LDAP is a protocol that provides access to directory services over a network. AD Query uses LDAP queries to retrieve information about users and groups from Active Directory Domain Controllers without installing any software on them. AD Query does not use Windows Management Instrumentation (WMI), Hypertext Transfer Protocol Secure (HTTPS), or Remote Desktop Protocol (RDP) to connect to Active Directory Domain Controllers. Reference: Check Point Certified Security Administrator (CCSA) R80.x Study Guide, Chapter 5: User Management and Authentication, page 69.

QUESTION 79
What is the main objective when using Application Control?

A. To filter out specific content.
B. To assist the firewall blade with handling traffic.
C. To see what users are doing.
D. Ensure security and privacy of information.

Correct Answer: D
Section:
Explanation:
The main objective when using Application Control is to ensure the security and privacy of information. Application Control is a security practice that blocks or restricts unauthorized applications from executing in ways that put data at risk. The control functions vary based on the business purpose of the specific application, but the main objective is to help ensure the privacy and security of data used by and transmitted between applications1. Application Control can also prevent malware, untrusted, or unwanted applications from running on the network, reducing the risks and costs associated with data breaches1. Application Control can also improve the overall network stability and performance by eliminating unnecessary or harmful applications1.
Application Control is not mainly used to filter out specific content, although it can be combined with other technologies such as URL filtering or content filtering to achieve that goal. Application Control is not mainly used to assist the firewall blade with handling traffic, although it can be integrated with firewall policies to enforce granular access rules based on applications. Application Control is not mainly used to see what users are doing, although it can provide visibility and reporting on application usage and activity.

QUESTION 80
A Linux administrator reviews a set of log output files and needs to identify files that contain any occurrence of the word denied. All log files containing entries in uppercase or lowercase letters should be included in the list. Which of the following commands should the administrator use to accomplish this task?

A. find . -type f -print | xrags grep -ln denied
B. find . -type f -print | xrags grep -nv denied
C. find . -type f -print | xrags grep -wL denied
D. find . -type f -print | xrags grep -li denied

Correct Answer: D
Section:
Explanation:
The command find . -type f -print | xargs grep -li denied will accomplish the task of identifying files that contain any occurrence of the word denied. The find command is a tool for searching for files and directories on Linux systems. The . is the starting point of the search, which means the current directory. The -type f option specifies the type of the file, which means regular file. The -print option prints the full file name on the standard output. The | is a pipe symbol that redirects the output of one command to the input of another command. The xargs command is a tool for building and executing commands from standard input. The grep command is a tool for searching for patterns in files or input. The -li option specifies the flags that the grep command should apply. The -l flag shows only the file names that match the pattern, instead of the matching lines. The -i flag ignores the case of the pattern, which means it matches both uppercase and lowercase letters. The denied is the pattern that the grep command should search for. The command find . -type f -print | xargs grep -li denied will find all the regular files in the current directory and its subdirectories, and then search for any occurrence of the word denied in those files, ignoring the case, and print only the file names that match the pattern. This will allow the administrator to identify files that contain any occurrence of the word denied. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not ignore the case of the pattern (find . -type f -print | xargs grep -ln denied or find . -type f -print | xargs grep -wL denied) or do not show the file names that match the pattern (find . -type f -print | xargs grep -nv denied). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Managing Logging and Monitoring, page 489.

QUESTION 81
A Linux administrator is installing a web server and needs to check whether web traffic has already been allowed through the firewall. Which of the following commands should the administrator use to accomplish this task?

A. firewalld query-service-http
B. firewall-cmd --check-service http
C. firewall-cmd --query-service http
D. firewalld --check-service http

Correct Answer: C
Section:
Explanation:
The command firewall-cmd --query-service http will accomplish the task of checking whether web traffic has already been allowed through the firewall. The firewall-cmd command is a tool for managing firewalld, which is a firewall service that provides dynamic and persistent network security on Linux systems. The firewalld uses zones and services to define the rules and policies for the network traffic. The zones are logical groups of network interfaces and sources that have the same level of trust and security. The services are predefined sets of ports and protocols that are associated with certain applications or functions. The --query-service http option queries whether a service is enabled in a zone. The http is the name of the service that the command should check. The http service represents the web traffic that uses the port 80 and the TCP protocol. The command firewallcmd --query-service http will check whether the http service is enabled in the default zone, which is usually the public zone. The command will return yes if the web traffic has already been allowed through the firewall, or no if the

web traffic has not been allowed through the firewall. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not exist (firewalld query-service-http or firewalld --check-service http) or do not query the service (firewall-cmd --check-service http instead of firewall-cmd --query-service http). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 392.

QUESTION 82
A systems administrator wants to permit access temporarily to an application running on port 1234/TCP on a Linux server. Which of the following commands will permit this traffic?

A. firewall-cmd ---new-service=1234/tcp
B. firewall-cmd ---service=1234 ---protocol=tcp
C. firewall-cmd ---add---port=1234/tcp
D. firewall-cmd ---add-whitelist-uid=1234

Correct Answer: C
Section:
Explanation:
Thefirewall-cmdcommand is used to manage firewalld, which is a firewall service for Linux systems that provides dynamic and persistent configuration of firewall rules. Firewalld uses zones and services to define different levels of trust and access for network connections.
To permit access temporarily to an application running on port 1234/TCP on a Linux server, the systems administrator can use thefirewall-cmd --add-port=1234/tcpcommand. This command will add a rule to the default zone (usually public) that allows incoming traffic on port 1234/TCP. The rule will only be effective until the next reload or restart of firewalld. To make the rule permanent, the administrator can add the--permanentoption to the command. The statement C is correct.
The statements A, B, and D are incorrect because they do not permit access to port 1234/TCP. Thefirewall-cmd --new-service=1234/tcpcommand does not exist. Thefirewall-cmd --service=1234 --protocol=tcpcommand does not work because 1234 is not a predefined service name in firewalld. Thefirewall-cmd --add-whitelist-uid=1234command does not exist.Reference: [How to Use FirewallD to Manage Firewall in Linux]

QUESTION 83
The development team wants to prevent a file from being modified by all users in a Linux system, including the root account. Which of the following commands can be used to accomplish this objective?

A. chmod / app/conf/file
B. setenforce / app/ conf/ file
C. chattr +i /app/conf/file
D. chmod 0000 /app/conf/file

Correct Answer: C
Section:
Explanation:
Thechattrcommand is used to change file attributes on Linux systems that support extended attributes, such as ext2, ext3, ext4, btrfs, xfs, and others. File attributes are flags that modify the behavior of files and directories.
To prevent a file from being modified by all users in a Linux system, including the root account, the development team can use thechattr +i /app/conf/filecommand. This command will set the immutable attribute (+i) on the file/app/conf/file, which means that the file cannot be deleted, renamed, linked, appended, or written to by any user or process. To remove the immutable attribute, the development team can use thechattr -i /app/conf/filecommand. The statement C is correct.
The statements A, B, and D are incorrect because they do not prevent the file from being modified by all users. Thechmod /app/conf/filecommand does not work because it requires an argument to specify the permissions to change. Thesetenforce /app/conf/filecommand does not work because it is used to change the SELinux mode, not file attributes. Thechmod 0000 /app/conf/filecommand will remove all permissions from the file, but it can still be modified by the root account.Reference: [How to Use chattr Command in Linux]

QUESTION 84
The development team wants to prevent a file from being modified by all users in a Linux system, including the root account. Which of the following commands can be used to accomplish this objective?

A. chmod / app/conf/file
B. setenforce / app/ conf/ file
C. chattr +i /app/conf/file

D. chmod 0000 /app/conf/file

Correct Answer: C
Section:
Explanation:
Thechattrcommand is used to change file attributes on Linux systems that support extended attributes, such as ext2, ext3, ext4, btrfs, xfs, and others. File attributes are flags that modify the behavior of files and directories. To prevent a file from being modified by all users in a Linux system, including the root account, the development team can use thechattr +i /app/conf/filecommand. This command will set the immutable attribute (+i) on the file/app/conf/file, which means that the file cannot be deleted, renamed, linked, appended, or written to by any user or process. To remove the immutable attribute, the development team can use thechattr -i /app/conf/filecommand. The statement C is correct.
The statements A, B, and D are incorrect because they do not prevent the file from being modified by all users. Thechmod /app/conf/filecommand does not work because it requires an argument to specify the permissions to change. Thesetenforce /app/conf/filecommand does not work because it is used to change the SELinux mode, not file attributes. Thechmod 0000 /app/conf/filecommand will remove all permissions from the file, but it can still be modified by the root account.Reference: [How to Use chattr Command in Linux]

QUESTION 85
An administrator would like to securely connect to a server and forward port 8080 on a local machine to port 80 on the server. Which of the following commands should the administrator use to satisfy both requirements?

A. ssh ---L 8080: localhost:80 admin@server
B. ssh ---R 8080: localhost:80 admin@server
C. ssh ---L 80 : localhost:8080 admin@server
D. ssh ---R 80 : localhost:8080 admin@server

Correct Answer: A
Section:
Explanation:
This command will create a local port forwarding, which means that connections from the SSH client are forwarded via the SSH server, then to a destination server. In this case, the destination server is the same as the SSH server (localhost), and the destination port is 80. The SSH client will listen on port 8080 on the local machine, and any connection to that port will be forwarded to port 80 on the server. This way, the administrator can securely access the web service running on port 80 on the server by using http://localhost:8080 on the local machine.
The other options are incorrect because:
B) ssh -R 8080:localhost:80 admin@server
This command will create a remote port forwarding, which means that connections from the SSH server are forwarded via the SSH client, then to a destination server. In this case, the destination server is the same as the SSH client (localhost), and the destination port is 80. The SSH server will listen on port 8080 on the remote machine, and any connection to that port will be forwarded to port 80 on the client. This is not what the administrator wants to do.
C) ssh -L 80:localhost:8080 admin@server
This command will also create a local port forwarding, but it will use port 80 on the local machine and port 8080 on the server. This is not what the administrator wants to do, and it may also fail if port 80 is already in use by another service on the local machine.
D) ssh -R admin@server
This command is incomplete and invalid. It does not specify any port numbers or destination addresses for the remote port forwarding. It will also fail if the SSH server does not allow remote port forwarding.
CompTIA Linux+ Certification Exam Objectives
How to Set up SSH Tunneling (Port Forwarding)

QUESTION 86
An engineer needs to insert a character at the end of the current line in the vi text editor. Which of the following will allow the engineer to complete this task?

A. p
B. r
C. bb
D. A
E. i

Correct Answer: D
Section:
Explanation:
The vi text editor is a popular and powerful tool for editing text files on Linux systems. The vi editor has two modes: command mode and insert mode. In command mode, the user can issue commands to manipulate the text, such as moving the cursor, deleting, copying, pasting, searching, replacing, and saving. In insert mode, the user can type text into the file. To switch from command mode to insert mode, the user can press various keys, such as i, a, o, I, A, or O. To switch from insert mode to command mode, the user can press the Esc key.
To insert a character at the end of the current line in the vi editor, the user can press the A key in command mode. This will move the cursor to the end of the line and switch to insert mode. Then, the user can type the desired character and press Esc to return to command mode. The statement D is correct.
The statements A, B, C, and E are incorrect because they do not perform the desired task. The p key in command mode will paste the previously copied or deleted text after the cursor. The r key in command mode will replace the character under the cursor with another character. The bb key in command mode will move the cursor back two words. The i key in command mode will switch to insert mode before the cursor.Reference: [How to Use vi Text Editor in Linux]

QUESTION 87
Which of the following specifications is used to perform disk encryption in a Linux system?

A. LUKS
B. TLS
C. SSL
D. NFS

Correct Answer: A
Section:
Explanation:
LUKS stands for Linux Unified Key Setup, which is a specification for disk encryption on Linux systems. LUKS allows users to encrypt partitions or entire disks using a passphrase or a key file. LUKS also supports multiple keys and key slots, which can be used to unlock the encrypted data. LUKS is compatible with various tools and utilities, such as cryptsetup, dm-crypt, and LVM.Reference: [How to Encrypt Partitions with LUKS on Linux]

QUESTION 88
As part of the requirements for installing a new application, the swappiness parameter needs to be changed to O. This change needs to persist across re-boots and be applied immediately. A Linux systems administrator is performing this change. Which of the following steps should the administrator complete to accomplish this task?

A. echo 'vm. swappiness---()' >> /etc/sysctl . conf && sysctl ---p
B. echo 'vrn. >> / proc/meminfo && sysctl ---a
C. sysctl ---v >> / proc/meminfo & & echo 'vm. swapiness=0'
D. sysctl ---h 'vm. swapiness---O' && echo / etc/vmswapiness

Correct Answer: A
Section:
Explanation:
To change the swappiness parameter to 0 and make it persistent across reboots and applied immediately, the administrator can perform the following steps:
Append the linevm.swappiness=0to the file/etc/sysctl.confusingecho 'vm.swappiness=0' >> /etc/sysctl.conf(A). This will set the swappiness parameter to 0 for future boots.
Reload the sysctl configuration usingsysctl -p(A). This will apply the changes to the current system without rebooting. The other commands will not achieve this task, but either write to a wrong file, use a wrong option, or have a syntax error.Reference:
[CompTIA Linux+ Study Guide], Chapter 8: Optimizing Linux Performance, Section: Tuning Kernel Parameters with sysctl
[How to Change Swappiness in Linux]

QUESTION 89
An administrator would like to mirror the website files on the primary web server, www1, to the backup web server, www2. Which of the following commands should the administrator use to most efficiently accomplish this task?

A. [wwwl] rsync ---a ---e ssh /var/www/html/ user1@www2 : /var/www/html

B. [wwwl] scp ---r /var/www/html user1@www2 : / var/www/html

C. [www2] cd /var/www/html; wget ---m http: //wwwl/

D. [wwwl] cd /var/www/html && tar cvf ---

Correct Answer: A
Section:
Explanation:
To mirror the website files on the primary web server, www1, to the backup web server, www2, the administrator can use the commandrsync -a -e ssh /var/www/html/ user1@www2:/var/www/html(A). This will synchronize all files and directories under/var/www/html/on www1 to/var/www/htmlon www2 using ssh as the remote shell. The-aoption will preserve all attributes and permissions of the files. The other commands will not mirror the website files, but either copy them once, download them from a web server, or archive them.Reference:
[CompTIA Linux+ Study Guide], Chapter 12: Troubleshooting Linux Systems, Section: Synchronizing Files with rsync
[How to Use rsync Command in Linux]

QUESTION 90
A Linux user is trying to execute commands with sudo but is receiving the following error:

```
$ sudo visudo
>>> /etc/sudoers: syntax error near line 28 <<<
sudo: parse error in /etc/sudoers near line 28
sudo: no valid sudoers sources found, quitting
```

The following output is provided:

```
# grep root /etc/shadow
root :* LOCK *: 14600 ::::::
```

Which of the following actions will resolve this issue?

A. Log in directly using the root account and comment out line 28 from /etc/sudoers.

B. Boot the system in single user mode and comment out line 28 from /etc/sudoers.

C. Comment out line 28 from /etc/sudoers and try to use sudo again.

D. Log in to the system using the other regular user, switch to root, and comment out line 28 from /etc/sudoers.

Correct Answer: B
Section:

QUESTION 91
A systems administrator is gathering information about a file type and the contents of a file. Which of the following commands should the administrator use to accomplish this task?

A. file filename

B. touch filename

C. grep filename

D. lsof filename

Correct Answer: A
Section:
Explanation:
The file command is used to determine the type of a file by examining its contents. It can recognize many different formats, such as text, binary, executable, compressed, image, audio, video, etc.It can also display some additional information about the file, such as encoding, size, dimensions, etc12

QUESTION 92

Due to performance issues on a server, a Linux administrator needs to termi-nate an unresponsive process. Which of the following commands should the administrator use to terminate the process immediately without waiting for a graceful shutdown?

A. kill -SIGKILL 5545
B. kill -SIGTERM 5545
C. kill -SIGHUP 5545
D. kill -SIGINT 5545

Correct Answer: A
Section:
Explanation:
To terminate an unresponsive process immediately without waiting for a graceful shutdown, the administrator can use the commandkill -SIGKILL 5545(A). This will send a signal to the process with the PID 5545 that cannot be ignored or handled by the process, and force it to stop. The other commands will send different signals that may allow the process to perform some cleanup or termination actions, or may be ignored by the process.Reference:
[CompTIA Linux+ Study Guide], Chapter 6: Managing Processes, Section: Killing Processes
[How to Kill Processes in Linux]

QUESTION 93
A systems administrator intends to use a UI-JID to mount a new partition per-manently on a Linux system. Which of the following commands can the adminis-trator run to obtain information about the UUIDs of all disks attached to a Linux system?

A. fcstat
B. blkid
C. dmsetup
D. lsscsi

Correct Answer: B
Section:
Explanation:
To obtain information about the UUIDs of all disks attached to a Linux system, the administrator can run the commandblkid(B). This will display the block device attributes, including the UUID, label, type, and partition information. The other commands are not related to this task.Reference:
[CompTIA Linux+ Study Guide], Chapter 5: Managing Filesystems and Logical Volumes, Section: Identifying Disks by UUID
[How to Use blkid Command in Linux]

QUESTION 94
A systems administrator creates a public key for authentication. Which of the following tools is most suitable to use when uploading the key to the remote servers?

A. scp
B. ssh-copy-id
C. ssh-agent
D. ssh-keyscan

Correct Answer: B
Section:
Explanation:
The best tool to use when uploading the public key to the remote servers is B. ssh-copy-id. This tool will copy the public key from the local computer to the remote server and append it to the authorized_keys file, which is used for public key authentication. This tool will also create the necessary directories and files on the remote server if they do not exist. The other tools are either not suitable or not relevant for this task. For example:
A) scp is a tool for securely copying files between hosts, but it does not automatically add the public key to the authorized_keys file.

C) ssh-agent is a tool for managing private keys and passphrases, but it does not upload the public key to the remote server.
D) ssh-keyscan is a tool for collecting public keys from remote hosts, but it does not upload the public key to the remote server.