CompTIA

Exam Questions XK0-005

CompTIA Linux+ Certification Exam

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1
A Linux administrator needs to remove software from the server. Which of the following RPM options should be used?

A. rpm -s
B. rm -d
C. rpm -q
D. rpm -e

Answer: D

Explanation:
The RPM option -e should be used to remove software from the server. The rpm command is a tool for managing software packages on RPM-based Linux distributions. The -e option stands for erase and removes the specified package from the system. This is the correct option to use to accomplish the task. The other options are incorrect because they either do not exist (-s or -d) or do not remove software (-q stands for query and displays information about the package).
References: CompTIA Linux+ (XK0-
005) Certification Study Guide, Chapter 16: Managing Software, page 489.

NEW QUESTION 2
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.

Answer: B

NEW QUESTION 3
A Linux administrator needs to obtain a list of all volumes that are part of a volume group. Which of the following commands should the administrator use to accomplish this task?

A. vgs
B. lvs
C. fdisk -1
D. pvs

Answer: B

Explanation:
The lvs command can be used to obtain a list of all volumes that are part of a volume group. This command will display information such as the name, size, attributes, and volume group of each logical volume in the system. The vgs command can be used to obtain a list of all volume groups in the system, not the volumes. The fdisk -1 command is invalid, as -1 is not a valid option for fdisk. The pvs command can be used to obtain a list of all physical volumes in the system, not the volumes. References: CompTIA Linux+ (XK0- 005) Certification Study Guide, Chapter 14: Managing Disk Storage, page 461.

NEW QUESTION 4
A user reported issues when trying to log in to a Linux server. The following outputs were received:
Given the outputs above. which of the following is the reason the user is una-ble to log in to the server?

A. User1 needs to set a long password.
B. User1 is in the incorrect group.
C. The user1 shell assignment incorrect.
D. The user1 password is expired.

Answer: D

Explanation:
The user1 password is expired. This can be inferred from the output of the chage -l user1 command, which shows the password expiration information for user1. The output shows that the password expired on 2020-10-01, and the account expired on 2020-10-08. This means that user1 cannot log in to the server unless the password and account are reactivated by the system administrator.
The other options are not correct based on the outputs above. User1 does not need to set a long password, because the output of the passwd -S user1 command shows that the password has a minimum length of 5 characters, which is met by user1's password. User1 is not in the incorrect group, because the output of the groups user1 command shows that user1 belongs to the app group, which is presumably the correct group for accessing the server. The user1 shell assignment is not incorrect, because the output of the grep user1
/etc/passwd command shows that user1 has /bin/bash as the default shell, which is a valid and common shell for Linux users.

NEW QUESTION 5
A Linux administrator wants to find out whether files from the wget package have been altered since they were installed. Which of the following commands will provide the correct information?

A. rpm -i wget
B. rpm -qf wget
C. rpm -F wget

D. rpm -V wget

Answer: D

Explanation:
The command that will provide the correct information about whether files from the wget package have been altered since they were installed is rpm -V wget. This command will use the rpm utility to verify an installed RPM package by comparing information about the installed files with information from the RPM database. The verification process can check various attributes of each file, such as size, mode, owner, group, checksum, capabilities, and so on. If any discrepancies are found, rpm will report them using a single letter code for each attribute.
The other options are not correct commands for verifying an installed RPM package. The rpm -i wget command is invalid because -i is used to install a package from a file, not to verify an installed package. The rpm -qf wget command will query which package owns wget as a file name or path name, but it will not verify its attributes. The rpm -F wget command will freshen (upgrade) an already installed package with wget as a file name or path name, but it will not verify its attributes.
References: rpm(8) - Linux manual
page; Using RPM to Verify Installed Packages

NEW QUESTION 6
A junior developer is unable to access an application server and receives the following output:

```
[root@server1 ~]# ssh dev2@172.16.25.126
dev2@172.16.25.126's password:
Permission denied, please try again.
dev2@172.16.25.126's password:
Permission denied, please try again.
dev2@172.16.25.126's password:
Account locked due to 4 failed logins
Account locked due to 5 failed logins
Last login: Mon Apr 22 21:21:06 2021 from 172.16.16.52
```

The systems administrator investigates the issue and receives the following output:

```
[root@server1 ~]# pam_tally2 --user=dev2
Login Failures Latest failure From
dev2 5 04/22/21 21:22:37 172.16.16.52
```

Which of the following commands will help unlock the account?

A. Pam_tally2 --user=dev2 —-quiet
B. pam_ tally2 --user=dev2
C. pam_tally2 -–user+dev2 —-quiet
D. pam_tally2 --user=dev2 —-reset

Answer: D

Explanation:
To unlock an account that has been locked due to login failures, the administrator can use the command pam_tally2 --user=dev2 --reset (D). This will reset the failure counter for the user “dev2” and allow the user to log in again. The other commands will not unlock the account, but either display or increase the failure count. References:
? [CompTIA Linux+ Study Guide], Chapter 4: Managing Users and Groups, Section: Locking Accounts with pam_tally2
? [How to Lock and Unlock User Account in Linux]

NEW QUESTION 7
An administrator needs to make an application change via a script that must be run only in console mode. Which of the following best represents the sequence the administrator should execute to accomplish this task?

A. systemct1 isolate multi-user.target sh script.shsystemct1 isolate graphical.target
B. systemct1 isolate graphical.target sh script.shsystemct1 isolate multi-user.target
C. sh script.shsystemct1 isolate multi-user.target systemct1 isolate graphical.target
D. systemct1 isolate multi-user.target systemct1 isolate graphical.targetsh script.sh

Answer: A

Explanation:
The correct answer is A. systemctl isolate multi-user.target sh script.sh systemctl isolate graphical.target
This sequence will allow the administrator to switch from the graphical mode to the console mode, run the script, and then switch back to the graphical mode.
The systemctl command is used to control the systemd system and service manager, which manages the boot targets and services on Linux systems. The isolate subcommand starts the unit specified on the command line and its dependencies and stops all others. The multi-user.target is a boot target that provides a text-based console login, while the graphical.target is a boot target that provides a graphical user interface. By using systemctl isolate, the administrator can change the boot target on the fly without rebooting the system.
The sh command is used to run a shell script, which is a file that contains a series of commands that can be executed by the shell. The script.sh is the name of the script that contains the application change that the administrator needs to make. By running sh script.sh, the administrator can execute the script in the console mode.
The other options are incorrect because:
* B. systemctl isolate graphical.target sh script.sh systemctl isolate multi-user.target
This sequence will switch from the console mode to the graphical mode, run the script, and then switch back to the console mode. This is not what the administrator wants to do, as the script must be run only in console mode.

* C. sh script.sh systemctl isolate multi-user.target systemctl isolate graphical.target
This sequence will run the script in the current mode, which may or may not be console mode, and then switch to console mode and back to graphical mode. This is not what the administrator wants to do, as the script must be run only in console mode.
* D. systemctl isolate multi-user.target systemctl isolate graphical.target sh script.sh
This sequence will switch from graphical mode to console mode and then back to graphical mode, without running the script at all. This is not what the administrator wants to do, as the script must be run only in console mode.
References:
? systemctl(1) - Linux manual page
? How to switch between the CLI and GUI on a Linux server
? How to PROPERLY boot into single user mode in RHEL/CentOS 7/8
? Changing Systemd Boot Target in Linux
? Exit Desktop to Terminal in Ubuntu 19.10

NEW QUESTION 8
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

Answer: C

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. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 321.

NEW QUESTION 9
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

Answer: B

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. References: Check Point Rugged Appliance Datasheet, page 1.

NEW QUESTION 10
A user is unable to remotely log on to a server using the server name server1 and port 22.
The Linux engineer troubleshoots the issue and gathers the following information: Which of the following is most likely causing the issue?

A. server 1 is not in the DNS.
B. sshd is running on a non-standard port.
C. sshd is not an active service.
D. serverl is using an incorrect IP address.

Answer: B

Explanation:
The sshd is the Secure Shell Daemon, which is a service that allows remote login to a Linux system using the SSH protocol. The output shows that the sshd is running on port 2222, which is a non-standard port for SSH. The default port for SSH is 22, which is what the user is trying to use. Therefore, the statement B is most likely causing the issue. The statements A, C, and D are incorrect because they do not explain why the user cannot log on using port 22. References: [How to Change SSH Port in Linux]

NEW QUESTION 10
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

Answer: C

Explanation:
The chattr command 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 the chattr +i /app/conf/file command. 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 the chattr -i /app/conf/file command. 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. The chmod /app/conf/file command does not work because it requires an argument to specify the permissions to change. The setenforce /app/conf/file command does not work because it is used to change the SELinux mode, not file attributes. The chmod 0000 /app/conf/file command will remove all permissions from the file, but it can still be modified by the root account. References: [How to Use chattr Command in Linux]

NEW QUESTION 13
A User on a Linux workstation needs to remotely start an application on a Linux server and then forward the graphical display of that application back to the Linux workstation. Which of the following would enable the user to perform this action?

A. ssh -X user@server application
B. ssh -y user@server application
C. ssh user@server application
D. ssh -D user@server application

Answer: A

Explanation:
The ssh -X option enables X11 forwarding, which allows the user to run graphical applications on the remote server and display them on the local workstation. The user needs to specify the username, the server address, and the application name after the ssh -X command. The remote server also needs to have X11Forwarding enabled and xauth installed for this to work. References:
? The web search result 8 explains how to run a GUI application through SSH by configuring both the SSH client and server.
? The web search result 6 provides a detailed answer on how to forward X over SSH to run graphics applications remotely, with examples and troubleshooting tips.
? The CompTIA Linux+ Certification Exam Objectives mention that the candidate should be able to “use SSH for remote access and management” as part of the System Operation and Maintenance domain1.

NEW QUESTION 14
The group owner of the / home/ test directory would like to preserve all group permissions on files created in the directory. Which of the following commands should the group owner execute?

A. chmod g+s /home/test
B. chgrp test /home/test
C. chmod 777 /home/test
D. chown —hR test /home/test

Answer: A

Explanation:
The correct answer is A. chmod g+s /home/test
This command will set the setgid bit on the /home/test directory, which means that any file or subdirectory created in the directory will inherit the group ownership of the directory. This way, the group permissions on files created in the directory will be preserved. The chmod command is used to change the permissions of files and directories. The g+s option is used to set the setgid bit for the group.
The other options are incorrect because:
* B. chgrp test /home/test
This command will change the group ownership of the /home/test directory to test, but it will not affect the group ownership of files created in the directory. The chgrp command is used to change the group of files and directories. The test /home/test arguments are used to specify the new group and the target directory.
* C. chmod 777 /home/test
This command will give read, write, and execute permissions to everyone (owner, group, and others) on the /home/test directory, but it will not affect the group ownership or permissions of files created in the directory. The chmod command is used to change the permissions of files and directories. The 777 argument is an octal number that represents the permissions in binary form.
* D. chown -hR test /home/test
This command will change the owner and group of the /home/test directory and all its contents recursively to test, but it will not preserve the original group permissions on files created in the directory. The chown command is used to change the owner and group of files and directories. The -hR option is used to affect symbolic links and operate on all files and directories recursively. The test /home/test arguments are used to specify the new owner and group and the target directory.
References:
? How to Set File Permissions Using chmod
? How to Use Chmod Command in Linux with Examples
? How to Use Chown Command in Linux with Examples
? [How to Use Chgrp Command in Linux with Examples]

NEW QUESTION 15
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

Answer: A

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). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 571.

NEW QUESTION 19
A systems administrator wants to delete app . conf from a Git repository. Which of the following commands will delete the file?

A. git tag ap
B. conf
C. git commit app . conf
D. git checkout app . conf
E. git rm ap
F. conf

Answer: D

Explanation:
To delete a file from a Git repository, the administrator can use the command git rm app.conf (D). This will remove the file “app.conf” from the working directory and stage it for deletion from the repository. The administrator can then commit the change with git commit -m "Delete app.conf" to finalize the deletion. The other commands will not delete the file, but either tag, commit, or checkout the file. References:
? [CompTIA Linux+ Study Guide], Chapter 10: Working with Git, Section: Deleting Files with Git
? [How to Delete Files from Git]

NEW QUESTION 20
A junior administrator is trying to set up a passwordless SSH connection to one of the servers. The administrator follows the instructions and puts the key in the authorized_key file at the server, but the administrator is still asked to provide a password during the connection.
Given the following output:

```
junior@server:-$ ls -lh .ssh/auth*
-rw------- 1 junior junior 566 sep 13 20:56 .ssh/authorized_key
```

Which of the following commands would resolve the issue and allow an SSH connection to be established without a password?

A. restorecon -rv .ssh/authorized_key
B. mv .ssh/authorized_key .ssh/authorized_keys
C. systemct1 restart sshd.service
D. chmod 600 mv .ssh/authorized_key

Answer: B

Explanation:
The command mv .ssh/authorized_key .ssh/authorized_keys will resolve the issue and allow an SSH connection to be established without a password. The issue is caused by the incorrect file name of the authorized key file on the server. The file should be named authorized_keys, not authorized_key. The mv command will rename the file and fix the issue. The other options are incorrect because they either do not affect the file name (restorecon or chmod) or do not restart the SSH service (systemct1). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 410.

NEW QUESTION 21
A Linux systems administrator receives reports from various users that an application hosted on a server has stopped responding at similar times for several days in a row. The administrator logs in to the system and obtains the following output:
Output 1:

```
[Tue Aug 31 16:36:42 2021] OOM: Kill process 43805 (java) score 249 or sacrifice child
[Tue Aug 31 16:36:42 2021] killed process 43805 (java) total-vm: 4446352kB, anon-rss: 4053140kB, file-rss: 68kB
```

Output 2:

```
Linux 3.10.0-328.13.1.x86_64 #1 (hostname) 31/08/2021 _x86_64_ (8 CPU)
16:00:01 PM     CPU    %user   %nice   %system   %iowait   %steal       %idle
16:10:01 PM     all    17.58    0.00      9.36      0.00     0.00       73.06
16:20:01 PM     all    22.34    0.00     11.75      0.00     0.00       65.91
16:30:01 PM     all    25.49    0.00     11.69      0.00        0       62.82
```

Output 3:

```
$ free -m
              total        used        free      shared  buff/cache   available
Mem:          16704       15026         174          92         619         793
Swap:             0           0           0
```

Which of the following should the administrator do to provide the BEST solution for the reported issue?

A. Configure memory allocation policies during business hours and prevent the Java process from going into a zombie state while the server is idle.
B. Configure a different nice value for the Java process to allow for more users and prevent the Java process from restarting during business hours.
C. Configure more CPU cores to allow for the server to allocate more processing and prevent the Java process from consuming all of the available resources.
D. Configure the swap space to allow for spikes in usage during peak hours and prevent the Java process from stopping due to a lack of memory.

Answer: D

Explanation:
Based on the output of the image sent by the user, the system requires more swap space to allow for spikes in usage during peak hours and prevent the Java process from stopping due to a lack of memory. The output shows that there is only 0 MB of swap space available on the system, which means that there is no room for swapping out memory pages when physical memory is full or low. The output also shows that there is only 793 MB of available memory on the system, which may not be enough to handle high- demand applications such as Java. This may cause Java to stop working due to insufficient memory or trigger an OutOfMemoryError exception. Configuring more swap space on the system would help to alleviate this issue by providing more virtual memory for applications and improving performance. Configuring memory allocation policies during business hours will not help to solve this issue, as it will not increase the amount of available

memory or swap space on the system. Configuring a different nice value for Java process will not help to solve this issue, as it will only affect its scheduling priority, not its memory consumption or allocation. Configuring more CPU cores will not help to solve this issue, as it will only increase processing power, not memory capacity or availability. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Memory and Process Execution, page 468.

NEW QUESTION 24
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

Answer: A

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). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 325.

NEW QUESTION 27
Users in the human resources department are trying to access files in a newly created directory. Which of the following commands will allow the users access to the files?

A. chattr
B. chgrp
C. chage
D. chcon

Answer: B

Explanation:
The chgrp command is used to change the group ownership of files and directories. By using this command, the administrator can assign the files in the newly created directory to the human resources group, which will allow the users in that group to access them. The other commands are not relevant for this task. For example:
? chattr is used to change the file attributes, such as making them immutable or append-only1.
? chage is used to change the password expiration information for a user account2.
? chcon is used to change the security context of files and directories, which is related to SELinux3.
References:
? The CompTIA Linux+ Certification Exam Objectives mention that the candidate should be able to “manage file and directory ownership and permissions” as part of the Hardware and System Configuration domain4.
? The web search result 2 explains how to use the chgrp command with examples.
? The web search result 3 compares the chmod and chgrp commands and their effects on file permissions.

NEW QUESTION 29
Which of the following tools is commonly used for creating CI/CD pipelines?

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

Answer: C

Explanation:
The tool that is commonly used for creating CI/CD pipelines is Jenkins. Jenkins is an open-source automation server that enables continuous integration and continuous delivery (CI/CD) of software projects. Jenkins allows developers to build, test, and deploy code changes automatically and frequently using various plugins and integrations. Jenkins also supports distributed builds, parallel execution, pipelines as code, and real-time feedback. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Source Code; Jenkins

NEW QUESTION 30
Following the migration from a disaster recovery site, a systems administrator wants a server to require a user to change credentials at initial login. Which of the following commands should
be used to ensure the aging attribute?

A. chage -d 2 user
B. chage -d 0 user
C. chage -E 0 user
D. chage -d 1 user

Answer: B

Explanation:
The chage command can be used to change the user password expiry information. The -d or --lastday option sets the last password change date. If the value is 0, the user will be forced to change the password at the next login. See chage command in Linux with examples and 10 chage command examples in Linux.

NEW QUESTION 31
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)

Answer: C

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. References: Check Point Certified Security Administrator (CCSA) R80.x Study Guide, Chapter 5: User Management and Authentication, page 69.

NEW QUESTION 32
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 && systemct1 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 && systemct1 restart systemd-journald.service

Answer: C

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 && systemct1 restart systemd-journald.service or cat /etc/systemd/journald.conf | awk '(print $1,$3)') or do not enable the Storage option (journalctl --list-boots && systemct1 restart systemd- journald.service). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Managing Logging and Monitoring, page 489.

NEW QUESTION 35
A systems administrator needs to check if the service systemd-resolved.service is running without any errors. Which of the following commands will show this information?

A. systemct1 status systemd-resolved.service
B. systemct1 enable systemd-resolved.service
C. systemct1 mask systemd-resolved.service
D. systemct1 show systemd-resolved.service

Answer: A

Explanation:
The command systemct1 status systemd-resolved.service will show the information about the service systemd-resolved.service. The systemct1 command is a tool for managing system services and units. The status option displays the current status of a unit, such as active, inactive, or failed. The output also shows the unit description, loaded configuration, process ID, memory usage, and recent log messages. This command will show if the service systemd-resolved.service is running without any errors. This is the
correct command to use to accomplish the task. The other options are incorrect because they either perform different actions (enable, mask, or show) or do not show the status of the service (systemct1 show systemd-resolved.service only shows the properties of the service, not the status). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, page 427.

NEW QUESTION 38
A Linux administrator needs to transfer a local file named accounts . pdf to a remote / tmp directory of a server with the IP address 10.10.10.80. Which of the following commands needs to be executed to transfer this file?

A. rsync user@10.10.10.80: /tmp accounts.pdf
B. scp accounts.pdf user@10.10.10.80:/tmp
C. cp user@10.10.10. 80: /tmp accounts.pdf
D. ssh accounts.pdf user@10.10.10.80: /tmp

Answer: B

Explanation:
The best command to use to transfer the local file accounts.pdf to the remote /tmp directory of the server with the IP address 10.10.10.80 is B. scp accounts.pdf user@10.10.10.80:/tmp. This command will use the secure copy protocol (scp) to copy the file from the local machine to the remote server over SSH. The command requires the username and password of the user on the remote server, as well as the full path of the destination directory.
The other commands are either incorrect or not suitable for this task. For example:
? A. rsync user@10.10.10.80:/tmp accounts.pdf will try to use the rsync command to synchronize files between the local and remote machines, but it has the wrong syntax and order of arguments. The source should come before the destination, and a colon (:) should separate the remote host and path.
? C. cp user@10.10.10.80:/tmp accounts.pdf will try to use the cp command to copy files, but it does not work over SSH and it has the wrong syntax and order of arguments. The source should come before the destination, and a colon (:) should separate the remote host and path.
? D. ssh accounts.pdf user@10.10.10.80:/tmp will try to use the ssh command to log into the remote server, but it has the wrong syntax and arguments. The username should come before the remote host, and a file name is not a valid argument for ssh.

NEW QUESTION 40
Which of the following will prevent non-root SSH access to a Linux server?

A. Creating the /etc/nologin file
B. Creating the /etc/nologin.allow file containing only a single line root
C. Creating the /etc/nologin/login.deny file containing a single line +all
D. Ensuring that /etc/pam.d/sshd includes account sufficient pam_nologin.so

Answer: A

Explanation:
This file prevents any non-root user from logging in to the system, regardless of the authentication method. The contents of the file are displayed to the user before the login is terminated. This can be useful for system maintenance or security reasons12.
References: 1: Creating the /etc/nologin File - Oracle 2: How to Restrict Log In Capabilities of Users on Ubuntu

NEW QUESTION 43
An administrator attempts to connect to a remote server by running the following command:
$ nmap 192.168.10.36
Starting Nmap 7.60 (https://nmap.org) at 2022-03-29 20:20 UTC Nmap scan report for www1 (192.168.10.36)
Host is up (0.000091s latency). Not shown: 979 closed ports PORT STATE SERVICE 21/tcp open ftp 22/tcp filtered ssh 631/tcp open ipp
Nmap done: 1 IP address (1 host up) scanned in 0.06 seconds
Which of the following can be said about the remote server?

A. A firewall is blocking access to the SSH server.
B. The SSH server is not running on the remote server.
C. The remote SSH server is using SSH protocol version 1.
D. The SSH host key on the remote server has expired.

Answer: A

Explanation:
This is because the port 22/tcp is shown as filtered by nmap, which means that nmap cannot determine whether the port is open or closed because a firewall or other device is blocking its probes. If the SSH server was not running on the remote server, the port would be shown as closed, which means that nmap received a TCP RST packet in response to its probe. If the remote SSH server was using SSH protocol version 1, the port would be shown as open, which means that nmap received a TCP SYN/ACK packet in response to its probe. If the SSH host key on the remote server had expired, the port would also be shown as open, but the SSH client would display a warning message about the host key verification failure. Therefore, the best explanation for the filtered state of the port 22/tcp is that a firewall is preventing nmap from reaching the SSH server.
You can find more information about nmap port states and how to interpret them in the following web search results:
? Nmap scan what does STATE=filtered mean?
? How to find ports marked as filtered by nmap
? Technical Tip: NMAP scan shows ports as filtered

NEW QUESTION 45
A systems administrator notices the process list on a mission-critical server has a large number of processes that are in state "Z" and marked as "defunct." Which of the following should the administrator do in an attempt to safely remove these entries from the process list?

A. Kill the process with PID 1.
B. Kill the PID of the processes.
C. Kill the parent PID of the processes.
D. Reboot the server.

Answer: C

Explanation:
As the web search results show, processes in state Z are defunct or zombie processes, which means they have terminated but their parent process has not reaped them properly. They do not consume any resources, but they occupy a slot in the process table. To remove them from the process list, the administrator needs to kill the parent process of the zombies, which will cause them to be reaped by the init process (PID 1). Killing the zombies themselves or the init process will not have any effect, as they are already dead. Rebooting the server may work, but it is not a safe or efficient option, as it may cause unnecessary downtime or data loss for a mission-critical server.
References
? Processes in a Zombie (Z) or Defunct State | Support | SUSE, paragraph 3
? linux - Zombie vs Defunct processes? - Stack Overflow, answer by admirableadmin
? How To Kill Zombie Processes on Linux | Linux Journal, paragraph 4

NEW QUESTION 49

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

Answer: DF

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. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, pages 408-410.

NEW QUESTION 53
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

Answer: D

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). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 381.

NEW QUESTION 57
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

Answer: B

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). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 567.

NEW QUESTION 61
A Linux administrator found many containers in an exited state. Which of the following commands will allow the administrator to clean up the containers in an exited state?

A. docker rm --all
B. docker rm $(docker ps -aq)
C. docker images prune *
D. docker rm --state exited

Answer: B

Explanation:
The command docker rm $(docker ps -aq) will allow the administrator to clean up the containers in an exited state. The docker command is a tool for managing Docker containers on Linux systems. Docker containers are isolated and lightweight environments that can run applications and services without affecting the host system. Docker uses images to create containers, which are files that contain the code, libraries, dependencies, and configuration of the applications and services. The rm option removes one or more containers. The $(docker ps -aq) is a command substitution that executes the command inside the parentheses and replaces it with the output. The docker ps - aq command lists all the containers, including the ones in an exited state, and shows only their IDs. The docker rm $(docker ps -aq) command will remove all the containers, including the ones in an exited state, by passing their IDs to the rm option. This will allow the administrator to clean up the containers in an exited state. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not exist (docker rm --all or docker rm --state exited) or do not remove the containers (docker images prune *). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 571.

NEW QUESTION 65
A Linux engineer needs to download a ZIP file and wants to set the nice of value to -10 for this new process. Which of the following commands will help to accomplish the task?

A. $ nice -v -10 wget https://foo.com/installation.zip
B. $ renice -v -10 wget https://foo.com/installation.2ip
C. $ renice -10 wget https://foo.com/installation.zip
D. $ nice -10 wget https://foo.com/installation.zip

Answer: D

Explanation:
The nice -10 wget https://foo.com/installation.zip command will help to accomplish the task of downloading a ZIP file and setting the nice value to -10 for this new process. The nice command can be used to run a program with a modified scheduling priority, which affects how much CPU time the process receives. The nice value ranges from -20 (highest priority) to 19 (lowest priority), and the default value is 0. The -10 option specifies the nice value to be used for the wget command, which will download the ZIP file from the given URL. The nice -v -10 wget https://foo.com/installation.zip command is incorrect, as -v is not a valid option for nice. The renice -v -10 wget https://foo.com/installation.zip command is incorrect, as renice is used to change the priority of an existing process, not a new one. The renice -10 wget https://foo.com/installation.zip command is incorrect for the same reason as above. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Memory and Process Execution, page 469.

NEW QUESTION 69
An administrator would like to list all current containers, regardless of their running state. Which of the following commands would allow the administrator to accomplish this task?

A. docker ps -a
B. docker list
C. docker image ls
D. docker inspect image

Answer: A

Explanation:
The best command to use to list all current containers, regardless of their running state, is A. docker ps -a. This command will show all containers, both running and stopped, with details such as container ID, image name, status, and ports. The other commands are either invalid or not relevant for this task. For example:
? B. docker list is not a valid command. There is no subcommand named list in docker.
? C. docker image ls will list all the images available on the local system, not the containers.
? D. docker inspect image will show detailed information about a specific image, not all the containers.

NEW QUESTION 74
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

Answer: A

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.
References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing the Linux Boot Process, page 267.

NEW QUESTION 79
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

Answer: D

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. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 559.

NEW QUESTION 81
A systems administrator needs to remove a disk from a Linux server. The disk size is 500G, and it is the only one that size on that machine. Which of the following commands can the administrator use to find the corresponding device name?

A. fdisk -V
B. partprobe -a
C. lsusb -t
D. lsscsi -s

Answer: D

Explanation:
The lsscsi command can list the SCSI devices on the system, along with their size and device name. The -s option shows the size of each device. The administrator can look for the device that has a size of 500G and note its device name. See lsscsi(8) - Linux man page and How to check Disk Interface Types in Linux.References1: https://linux.die.net/man/8/lsscsi2: https://www.golinuxcloud.com/check-disk-type-linux/

NEW QUESTION 86
A systems administrator is trying to track down a rogue process that has a TCP listener on a network interface for remote command-and-control instructions. Which of the following commands should the systems administrator use to generate a list of rogue process names? (Select two).

A. netstat -antp | grep LISTEN
B. lsof -iTCP | grep LISTEN
C. lsof -i:22 | grep TCP
D. netstat -a | grep TCP
E. nmap -p1-65535 | grep -i tcp
F. nmap -sS 0.0.0.0/0

Answer: AB

Explanation:
The best commands to use to generate a list of rogue process names that have a TCP listener on a network interface are A. netstat -antp | grep LISTEN and B. lsof -iTCP | grep LISTEN. These commands will show the process ID (PID) and name of the processes that are listening on TCP ports, which can be used to identify any suspicious or unauthorized processes. The other commands are either not specific enough, not valid, or not relevant for this task. For example:
? C. lsof -i:22 | grep TCP will only show the processes that are listening on port 22, which is typically used for SSH, and not any other ports.
? D. netstat -a | grep TCP will show all the TCP connections, both active and listening, but not the process names or IDs.
? E. nmap -p1-65535 | grep -i tcp will scan all the TCP ports on the local host, but not show the process names or IDs.
? F. nmap -sS 0.0.0.0/0 will perform a stealth scan on the entire internet, which is not only impractical, but also illegal in some countries.

NEW QUESTION 89
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

Answer: D

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. References: How to Use Public Key Authentication with SSH; chmod(1) - Linux manual

page

NEW QUESTION 93
A Linux administrator is troubleshooting an issue in which users are not able to access https://portal.comptia.org from a specific workstation. The administrator runs a few commands and receives the following output:

```
# cat /etc/hosts
10.10.10.55 portal.comptia.org

# host portal.comptia.org
portal.comptia.org has address 192.168.1.55

#cat /etc/resolv.conf
nameserver 10.10.10.5
```

Which of the following tasks should the administrator perform to resolve this issue?

A. Update the name server in resol
B. conf to use an external DNS server.
C. Remove the entry for portal . comptia.org from the local hosts file.
D. Add a network route from the 10.10.10.0/24 to the 192.168.0.0/16.
E. Clear the local DNS cache on the workstation and rerun the host command.

Answer: B

Explanation:
The best task to perform to resolve this issue is B. Remove the entry for portal.comptia.org from the local hosts file. This is because the local hosts file has a wrong entry that maps portal.comptia.org to 10.10.10.55, which is different from the actual IP address of 192.168.1.55 that is returned by the DNS server. This causes a mismatch and prevents the workstation from accessing the website. By removing or correcting the entry in the hosts file, the workstation will use the DNS server to resolve the domain name and access the website successfully.
To remove or edit the entry in the hosts file, you need to have root privileges and use a text editor such as vi or nano. For example, you can run the command:
sudo vi /etc/hosts
and delete or modify the line that says: 10.10.10.55 portal.comptia.org
Then save and exit the file.

NEW QUESTION 94
A systems administrator is installing various software packages using a pack-age manager. Which of the following commands would the administrator use on the Linux server to install the package?

A. winget
B. softwareupdate
C. yum-config
D. apt

Answer: D

NEW QUESTION 98
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

Answer: C

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 |). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, page 351.

NEW QUESTION 102
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

Answer: D

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 owner-group-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). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing File Permissions and Ownership, page 352.

NEW QUESTION 107
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. systemct1 start ipv4_forwarding

Answer: B

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 systemct1), or the wrong option (--enable or start). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 378.

NEW QUESTION 108
A Linux administrator needs to expand a volume group using a new disk. Which of the following options presents the correct sequence of commands to accomplish the task?

A. partprobe vgcreate lvextend
B. lvcreate fdisk partprobe
C. fdisk partprobe mkfs
D. fdisk pvcreate vgextend

Answer: D

Explanation:
The correct sequence of commands to expand a volume group using a new disk is fdisk, pvcreate, vgextend. The fdisk command can be used to create a partition on the new disk with the type 8e (Linux LVM). The pvcreate command can be used to initialize the partition as a physical volume for LVM. The vgextend command can be used to add the physical volume to an existing volume group. The partprobe command can be used to inform the kernel about partition table changes, but it is not necessary in this case. The vgcreate command can be used to create a new volume group, not expand an existing one. The lvextend command can be used to extend a logical volume, not a volume group. The lvcreate command can be used to create a new logical volume, not expand a volume group. The mkfs command can be used to create a filesystem on a partition or a logical volume, not expand a volume group. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Disk Storage, pages 462-463.

NEW QUESTION 111
A systems administrator is implementing a new service task with systems at startup and needs to execute a script entitled test.sh with the following content:

```
TIMESTAMP=$ (date '+%Y-%m-%d %H:%M:%S')
echo "helpme.service: timestamp $(Timestamp)" | systemd-cat -p info
sleep 60
done
```

The administrator tries to run the script after making it executable with chmod +x; however, the script will not run. Which of the following should the administrator do to address this issue? (Choose two.)

A. Add #!/bin/bash to the bottom of the script.
B. Create a unit file for the new service in /etc/systemd/system/ with the name helpme.service in the location.
C. Add #!//bin/bash to the top of the script.
D. Restart the computer to enable the new service.
E. Create a unit file for the new service in /etc/init.d with the name helpme.service in the location.
F. Shut down the computer to enable the new service.

Answer: BC

Explanation:
The administrator should do the following two things to address the issue:
? Add #!/bin/bash to the top of the script. This is called a shebang line and it tells the system which interpreter to use to execute the script. Without this line, the script will not run properly. The shebang line should be the first line of the script and should start with #! followed by the path to the interpreter. In this case, the interpreter is bash and the path is /bin/bash. The other option (A) is incorrect because the shebang line should be at the top, not the bottom of the script.
? Create a unit file for the new service in /etc/systemd/system/ with the name helpme.service in the location. This is necessary to register the script as a systemd service and enable it to run at startup. A unit file is a configuration file that defines the properties and behavior of a service, such as the description, dependencies, start and stop commands, and environment variables. The unit file should have the extension .service and should be placed in the /etc/systemd/system/ directory. The other option (E) is incorrect because /etc/init.d is the directory for init scripts, not systemd services.
References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, pages 429-430.

NEW QUESTION 116
An administrator added the port 2222 for the SSH server on myhost and restarted the SSH server. The administrator noticed issues during the startup of the service. Given the following outputs:

```
$ ssh -p 2222 myhost
ssh:connect to host myhost on port 2222: Connection refused

$ nmap -p 2222 myhost
Starting Nmap 7.70 ( https://nmap.org ) at 2022-10-17 21:12 EEST
Nmap scan report for myhost (10.7.3.26)
Host is up (0.00027s latency).
rDNS record for 10.7.3.26: myhost
 PORT     STATE  SERVICE
2222/tcp closed EtherNetIP-1
MAC Address: 52:54:00:F5:DF:F8 (QEMU virtual NIC)
 Nmap done: 1 IP address (1 host up) scanned in 0.57 seconds

$ systemctl status sshd
     • sshd.service - OpenSSH server daemon
   Loaded: loaded (/usr/lib/systemd/system/sshd.service; enabled; vendor preset: enabled)
   Active: active (running) since Mon 2022-10-17 19:40:07 CEST; 36min ago
     Docs: man:sshd(8)
           man:sshd_config(5)
Main PID: 13186 (sshd)
    Tasks: 1 (limit: 12373)
   Memory: 1.1M
   CGroup: /system.slice/sshd.service
           └─13186 /usr/sbin/sshd -D -oCiphers=aes256-gcm@openssh.com

Oct 17 19:40:07 myhost systemd[1]: Starting OpenSSH server daemon...
Oct 17 19:40:07 myhost sshd[13186]: error: Bind to port 2222 on 0.0.0.0 failed: Permission denied.
Oct 17 19:40:07 myhost systemd[1]: Started OpenSSH server daemon.
Oct 17 19:40:07 myhost sshd[13186]: Server listening on 0.0.0.0 port 22.
```

Which of the following commands will fix the issue?

A. semanage port -a -t ssh_port_t -p tcp 2222
B. chcon system_u:object_r:ssh_home_t /etc/ssh/*
C. iptables -A INPUT -p tcp -- dport 2222 -j ACCEPT
D. firewall-cmd -- zone=public -- add-port=2222/tcp

Answer: A

Explanation:
The correct answer is A. semanage port -a -t ssh_port_t -p tcp 2222
This command will allow the SSH server to bind to port 2222 by adding it to the SELinux policy. The semanage command is a utility for managing SELinux policies. The port subcommand is used to manage network port definitions. The -a option is used to add a new record, the -t option is used to specify the SELinux type, the -p option is used to specify the protocol, and the tcp 2222 argument is used to specify the port number. The ssh_port_t type is the default type for SSH ports in SELinux.
The other options are incorrect because:
* B. chcon system_u:object_r:ssh_home_t /etc/ssh/*
This command will change the SELinux context of all files under /etc/ssh/ to system_u:object_r:ssh_home_t, which is not correct. The ssh_home_t type is used for user home directories that are accessed by SSH, not for SSH configuration files. The correct type for SSH configuration files is sshd_config_t.
* C. iptables -A INPUT -p tcp --dport 2222 -j ACCEPT
This command will add a rule to the iptables firewall to accept incoming TCP connections on port 2222. However, this is not enough to fix the issue, as SELinux will still block the SSH server from binding to that port. Moreover, iptables may not be the default firewall service on some Linux distributions, such as Fedora or CentOS, which use firewalld instead.
* D. firewall-cmd --zone=public --add-port=2222/tcp
This command will add a rule to the firewalld firewall to allow incoming TCP connections on port 2222 in the public zone. However, this is not enough to fix the issue, as SELinux will still block the SSH server from binding to that port. Moreover, firewalld may not be installed or enabled on some Linux distributions, such as Ubuntu or Debian, which use iptables instead.
References:
? How to configure SSH to use a non-standard port with SELinux set to enforcing
? Change SSH Port on CentOS/RHEL/Fedora With SELinux Enforcing
? How to change SSH port when SELinux policy is enabled

NEW QUESTION 121
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.

Answer: CF

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). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing the Linux Boot Process, pages 267-268.

NEW QUESTION 122
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,
"systemct1 reboot" to reboot, "systemct1 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

Answer: C

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). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 321.

NEW QUESTION 127
A Linux administrator is creating a new sudo profile for the accounting user. Which of the following should be added by the administrator to the sudo configuration file so that the accounting user can run / opt/ acc/ report as root?

A. accounting localhost=/opt/acc/report
B. accounting ALL=/opt/acc/report
C. %accounting ALL=(ALL) NOPASSWD: /opt/acc/report
D. accounting /opt/acc/report= (ALL) NOPASSWD: ALL

Answer: C

Explanation:
This answer allows the accounting user to run the /opt/acc/report command as root on any host without entering a password. The % sign indicates that accounting is a group name, not a user name. The ALL keyword means any host, any user, and any command, depending on the context. The NOPASSWD tag overrides the default behavior of sudo, which is to ask for the user's password.
The other answers are incorrect for the following reasons:
? A. accounting localhost=/opt/acc/report
? B. accounting ALL=/opt/acc/report
? D. accounting /opt/acc/report= (ALL) NOPASSWD: ALL

NEW QUESTION 132
A Linux administrator needs to ensure that Java 7 and Java 8 are both locally available for developers to use when deploying containers. Currently only Java 8 is available. Which of the following commands should the administrator run to ensure both versions are available?

A. docker image load java:7
B. docker image pull java:7
C. docker image import java:7
D. docker image build java:7

Answer: B

Explanation:
The command that the administrator should run to ensure that both Java 7 and Java 8 are locally available for developers to use when deploying containers is docker image pull java:7. This command will use the docker image pull subcommand to download the java:7 image from Docker Hub, which is the default registry for Docker images. The java:7 image contains Java 7 installed on a Debian-based Linux system. The administrator can also specify a different registry by using the syntax registry/repository:tag.
The other options are not correct commands for ensuring that both Java 7 and Java 8 are locally available for developers to use when deploying containers. The docker image load java:7 command will load an image from a tar archive or STDIN, not from a registry. The docker image import java:7 command will create a new filesystem image from the contents of a tarball, not from a registry. The docker image build java:7 command will build an image from a Dockerfile, not from a registry. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks; docker image pull | Docker Docs

NEW QUESTION 135
A systems administrator is tasked with creating an Ansible playbook to automate the installation of patches on several Linux systems. In which of the following languages should the playbook be written?

A. SQL
B. YAML
C. HTML
D. JSON

Answer: B

Explanation:
The language that the playbook should be written in is YAML. YAML stands for YAML Ain't Markup Language, which is a human-readable data serialization language. YAML is commonly used for configuration files and data exchange. YAML uses indentation, colons, dashes, and brackets to represent the structure and values of the data. YAML also supports comments, variables, expressions, and functions. Ansible is an open-source tool for automating tasks and managing configuration on Linux systems. Ansible uses YAML to write playbooks, which are files that define the desired state and actions for the systems. Playbooks can be used to automate the installation of patches on several Linux systems by specifying the hosts, tasks, modules, and parameters. The language that the playbook should be written in is YAML. This is the correct answer to the question. The other options are incorrect because they are not the languages that Ansible uses for playbooks (SQL, HTML, or JSON). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 549.

NEW QUESTION 139
A systems administrator has been tasked with disabling the nginx service from the environment to prevent it from being automatically and manually started. Which of the following commands will accomplish this task?

A. systemct1 cancel nginx
B. systemct1 disable nginx
C. systemct1 mask nginx
D. systemct1 stop nginx

Answer: C

Explanation:
The command systemct1 mask nginx disables the nginx service from the environment and prevents it from being automatically and manually started. This command creates a symbolic link from the service unit file to /dev/null, which makes the service impossible to start. This is the correct way to accomplish the task. The other options are incorrect because they either do not exist (systemct1 cancel nginx), do not prevent manual start (systemct1 disable nginx), or do not prevent automatic start (systemct1 stop nginx). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, page 429.

NEW QUESTION 140
A systems administrator is configuring a Linux system so the network traffic from the internal network 172.17.0.0/16 going out through the eth0 interface would appear as if it was sent directly from this interface. Which of the following commands will accomplish this task?

A. iptables -A POSTROUTING -s 172.17.0.0/16 -o eth0 -j MASQUERADE
B. firewalld -A OUTPUT -s 172.17.0.0/16 -o eth0 -j DIRECT
C. nmcli masq-traffic eth0 -s 172.17.0.0/16 -j MASQUERADE
D. ifconfig -- nat eth0 -s 172.17.0.0/16 -j DIRECT

Answer: A

Explanation:
This command will use the iptables tool to append a rule to the POSTROUTING chain of the nat table, which will match any packet with a source address of 172.17.0.0/16 and an output interface of eth0, and apply the MASQUERADE target to it. This means that the packet will have its source address changed to the address of the eth0 interface, effectively hiding the internal network behind a NAT12.
References: 1: Iptables NAT and Masquerade rules - what do they do? 2: Routing from docker containers using a different physical network interface and default gateway

NEW QUESTION 144
A Linux administrator needs to resolve a service that has failed to start. The administrator runs the following command:

```
ls -l startup file
```

The following output is returned

```
----------. root root 81k Sep 13 19:01 startupfile
```

Which of the following is MOST likely the issue?

A. The service does not have permissions to read write the startupfile.
B. The service startupfile size cannot be 81k.
C. The service startupfile cannot be owned by root.
D. The service startupfile should not be owned by the root group.

Answer: A

Explanation:
The most likely issue is that the service does not have permissions to read or write the startupfile. The output of systemct1 status startup.service shows that the service has failed to start and the error message is “Permission denied”. The output of ls -l /etc/startupfile shows that the file has the permissions -rw-r--r--, which means that only the owner (root) can read and write the file, while the group (root) and others can only read the file. The service may not run as root and may need write access to the file. The administrator should change the permissions of the file by using the chmod command and grant write access to the group or others, or change the owner or group of the file by using the chown command and assign it to the user or group that runs the service. The other options are incorrect because they are not supported by the outputs. The file size, owner, and group are not the causes of the issue. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, pages 345-346.

NEW QUESTION 149
A new application container was built with an incorrect version number. Which of the following commands should be used to rename the image to match the correct version 2.1.2?

A. docker tag comptia/app:2.1.1 comptia/app:2.1.2
B. docker push comptia/app:2.1.1 comptia/app:2.1.2
C. docker rmi comptia/app:2.1.1 comptia/app:2.1.2
D. docker update comptia/app:2.1.1 comptia/app:2.1.2

Answer: A

Explanation:
The best command to use to rename the image to match the correct version 2.1.2 is A. docker tag comptia/app:2.1.1 comptia/app:2.1.2. This command will create a new tag for the existing image with the new version number, without changing the image content or ID. The other commands are either incorrect or not suitable for this task. For example:
? B. docker push comptia/app:2.1.1 comptia/app:2.1.2 will try to push two images to a remote repository, but it does not rename the image locally.
? C. docker rmi comptia/app:2.1.1 comptia/app:2.1.2 will try to remove two images
from the local system, but it does not rename the image.
? D. docker update comptia/app:2.1.1 comptia/app:2.1.2 will try to update the configuration of a running container, but it does not rename the image.

NEW QUESTION 150
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.

Answer: D

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. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Networking

NEW QUESTION 151
A Linux administrator is scheduling a system job that runs a script to check available disk space every hour. The Linux administrator does not want users to be able to start the job. Given the following:

```
[Unit]
Description=Check available disk space
RefuseManualstart=yes
RefuseManualStop=yes

[Timer]
Persistent=true
OnCalendar=*-*-*-*:00:00
Unit=checkdiskspace.service

[Install]
WantedBy=timers.target
```

The Linux administrator attempts to start the timer service but receives the following error message:

```
Failed to start checkdiskspace.timer: Operation refused ...
```

Which of the following is MOST likely the reason the timer will not start?

A. The checkdiskspace.timer unit should be enabled via systemct1.
B. The timers.target should be reloaded to get the new configuration.
C. The checkdiskspace.timer should be configured to allow manual starts.
D. The checkdiskspace.timer should be started using the sudo command.

Answer: C

Explanation:
The most likely reason the timer will not start is that the checkdiskspace.timer should be configured to allow manual starts. By default, systemd timers do not allow manual activation via systemct1 start, unless they have RefuseManualStart=no in their [Unit] section. This option prevents users from accidentally starting timers that are meant to be controlled by other mechanisms, such as calendar events or dependencies. To enable manual starts for checkdiskspace.timer, the administrator should add RefuseManualStart=no to its [Unit] section and reload systemd. The other options are not correct reasons for the timer not starting. The checkdiskspace.timer unit does not need to be enabled via systemct1 enable, because enabling a timer only makes it start automatically at boot time or after a system reload, but
does not affect manual activation. The timers.target does not need to be reloaded to get the new configuration, because reloading a target only affects units that have a dependency on it, but does not affect manual activation. The checkdiskspace.timer does not need to be started using the sudo command, because the administrator is already running systemct1 as root, as indicated by the # prompt. References: systemd.timer(5) - Linux manual page; systemct1(1) - Linux manual page

NEW QUESTION 154
A Linux administrator needs to correct the permissions of a log file on the server. Which of the following commands should be used to set filename.log permissions to -rwxr—r--. ?

A. chmod 755 filename.log
B. chmod 640 filename.log
C. chmod 740 filename.log
D. chmod 744 filename.log

Answer: A

Explanation:
The command chmod 755 filename.log should be used to set filename.log permissions to -rwxr--r--. The chmod command is a tool for changing file permissions on Linux file systems. The permissions can be specified in octal notation, where each digit represents the permissions for the owner, group, and others respectively. The permissions are encoded as follows:
? 0: no permission
? 1: execute permission
? 2: write permission
? 4: read permission
? 5: read and execute permissions (4 + 1)
? 6: read and write permissions (4 + 2)
? 7: read, write, and execute permissions (4 + 2 + 1)
The command chmod 755 filename.log will set the permissions to -rwxr--r--, which means that the owner has read, write, and execute permissions (7), the group has read and execute permissions (5), and others have read and execute permissions (5). This is the correct command to use to accomplish the task. The other options are incorrect because they either set the wrong permissions (chmod 640, chmod 740, or chmod 744) or do not exist (chmod -G). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, page 345.

NEW QUESTION 155
After installing a new version of a package, a systems administrator notices a new version of the corresponding, service file was Installed In order to use the new version of the, service file, which of the following commands must be Issued FIRST?

A. systemct1 status
B. systemct1 stop
C. systemct1 reinstall
D. systemct1 daemon-reload

Answer: D

Explanation:
After installing a new version of a package that includes a new version of the corresponding service file, the systemct1 daemon-reload command must be issued first in order to use the new version of the service file. This command will reload the systemd manager configuration and read all unit files that have changed on disk. This will ensure that systemd recognizes the new service file and applies its settings correctly. The systemct1 status command will display information about a service unit, but it will not reload the configuration. The systemct1 stop command will stop a service unit, but it will not reload the configuration. The systemct1 reinstall command does not exist. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: System Maintenance and Operation, page 518.

NEW QUESTION 159
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

Answer: D

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. References: CompTIA Linux+ (XK0- 005) Certification Study Guide, Chapter 13: Networking Fundamentals, page 436.

NEW QUESTION 163
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

Answer: B

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). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 381.

NEW QUESTION 168
A Linux systems administrator receives a notification that one of the server's filesystems is full. Which of the following commands would help the administrator to identify this filesystem?

A. lsblk
B. fdisk
C. df -h
D. du -ah

Answer: C

Explanation:
The df -h command can be used to identify the filesystem that is full. This command displays the disk usage of each mounted filesystem in a human-readable format, showing the total size, used space, available space, and percentage of each filesystem. The lsblk command displays information about block devices, not filesystems. The fdisk command can be used to manipulate partition tables, not check disk usage. The du -ah command displays the disk usage of each file and directory in a human-readable format, not the filesystems. References: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 14: Managing Disk Storage, page 454.

NEW QUESTION 169
After connecting to a remote host via SSH, an administrator attempts to run an application but receives the following error:
[user@workstation ~]$ ssh admin@srv1 Last login: Tue Mar 29 18:03:34 2022
[admin@srvl ~] $ /usr/local/bin/config_manager Error: cannot open display:
[admin@srv1 ~] $
Which of the following should the administrator do to resolve this error?

A. Disconnect from the SSH session and reconnect using the ssh -x command.
B. Add Options X11 to the /home/admin/.ssh/authorized_keys file.
C. Open port 6000 on the workstation and restart the firewalld service.
D. Enable X11 forwarding in /etc/ssh/ssh_config and restart the server.

Answer: A

Explanation:
The error indicates that the application requires an X11 display, but the SSH session does not forward the X11 connection. To enable X11 forwarding, the administrator needs to use the ssh -X option, which requests X11 forwarding with authentication spoofing. This will set the DISPLAY environment variable on the

remote host and allow the application to open a window on the local display.
References
? CompTIA Linux+ (XK0-005) Certification Study Guide, page 314
? Open a window on a remote X display (why “Cannot open display”)?, answer by Gilles ‘SO- stop being evil’

NEW QUESTION 170
A systems engineer is adding a new 1GB XFS filesystem that should be temporarily mounted under /ops/app. Which of the following is the correct list of commands to achieve this goal?

A.
```
pvcreate -L1G /dev/app
mkfs.xfs /dev/app
mount /dev/app /opt/app
```

B.
```
parted /dev/sdb --script mkpart primary xfs 1GB
mkfs.xfs /dev/sdb
mount /dev/sdb /opt/app
```

C.
```
lvs --create 1G --name app
mkfs.xfs /dev/app
mount /dev/app /opt/app
```

D.
```
lvcreate -L 1G -n app app_vq
mkfs.xfs /dev/app_vg/app
mount /dev/app_vg/app /opt/app
```

Answer: D

Explanation:
The list of commands in option D is the correct way to achieve the goal. The commands are as follows:
? fallocate -l 1G /ops/app.img creates a 1GB file named app.img under the /ops directory.
? mkfs.xfs /ops/app.img formats the file as an XFS filesystem.
? mount -o loop /ops/app.img /ops/app mounts the file as a loop device under the /ops/app directory. The other options are incorrect because they either use the wrong commands (dd or truncate instead of fallocate), the wrong options (-t or - f instead of -o), or the wrong order of arguments (/ops/app.img /ops/app instead of /ops/app /ops/app.img). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 323-324.

NEW QUESTION 175
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

Answer: A

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). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 9: Working with the Linux Shell, page 295.

NEW QUESTION 179
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

Answer: D

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. References: Handling package management history; dnf(8) - Linux manual page

NEW QUESTION 181
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.

Answer: B

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. References: CompTIA Linux+ Study Guide, Fourth Edition, page 457-458, 461-462.

NEW QUESTION 186
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.

Answer: C

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 systemct1 status checkdiskspace.service. The checkdiskspace.service is enabled, as shown by the output of systemct1 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. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, page 437.

NEW QUESTION 187
An administrator attempts to rename a file on a server but receives the following error.

```
mv: cannot move 'files/readme.txt' to 'files/readme.txt.orig': Operation not permitted.
```

The administrator then runs a few commands and obtains the following output:

```
$ ls -ld files/
 drwxrwxrwt.1    users    users    20    Sep 10 15:15    files/
$ ls -a files/
 drwxrwxrwt.1    users    users    20    Sep 10 15:15    .
 drwxr-xr-x.1    users    users    32    Sep 10 15:15    ..
 -rw-rw-r--.1    users    users     4    Sep 12 10:34    readme.txt
```

Which of the following commands should the administrator run NEXT to allow the file to be renamed by any user?

A. chgrp reet files
B. chacl -R 644 files
C. chown users files
D. chmod -t files

Answer: D

Explanation:
The command that the administrator should run NEXT to allow the file to be renamed by any user is chmod -t files. This command uses the chmod tool, which is used to change file permissions and access modes. The -t option removes (or sets) the sticky bit on a directory, which restricts deletion or renaming of files within that directory to only their owners or root. In this case, since files is a directory with sticky bit set (indicated by t in drwxrwxrwt), removing it will allow any user to rename or delete files within that directory. The other options are not correct commands for allowing any user to rename files within files directory. The chgrp reet files command will change the group ownership of files directory to reet, but it will not affect its permissions or access modes. The chacl -R 644 files command is invalid, as chacl is used to change file access control lists (ACLs), not permissions or access modes. The chown users files command will change the user ownership of files directory to users, but it will not affect its permissions or access modes. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing Users and Groups; chmod(1) - Linux manual page

NEW QUESTION 188
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/ director
B. systemct1 enable cleanupsystemct1 is-enabled cleanup
C. Create a unit file in the /etc/ske1/ director
D. systemct1 enable cleanupsystemct1 is-enabled cleanup
E. Create a unit file in the /etc/systemd/system/ director
F. systemct1 enable cleanupsystemct1 is-enabled cleanup
G. Create a unit file in the /etc/sysctl.d/ director
H. systemct1 enable cleanupsystemct1 is-enabled cleanup

Answer: C

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:
? Run the command systemct1 enable cleanup. This command will enable the service and create the necessary symbolic links to start the service at boot.
? Run the command systemct1 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. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing System Services, pages 457-459.

NEW QUESTION 189
A junior administrator updated the PostgreSQL service unit file per the data-base administrator's recommendation. The service has been restarted, but changes have not been applied. Which of the following should the administrator run for the changes to take effect?

A. Systemct1 get—default
B. systemct1 daemon—reload
C. systemct1 enable postgresq1
D. systemct1 mask postgresq1

Answer: B

Explanation:

To apply changes to a systemd service unit file, the administrator needs to reload the systemd daemon using the command systemct1 daemon-reload (B). This will make systemd aware of the new or changed unit files. The other commands will not reload the systemd daemon or apply the changes. References:
? [CompTIA Linux+ Study Guide], Chapter 7: Managing System Services, Section:
Modifying Systemd Services
? [How to Reload Systemd Services]

NEW QUESTION 192
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:

```
drwx------. 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.

Answer: D

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-xr-x), 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. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 413.

NEW QUESTION 193
A Linux systems administrator needs to persistently enable IPv4 forwarding in one of the Linux systems. Which of the following commands can be used together to accomplish this task? (Choose two.)

A. sysctl net.ipv4.ip_forward
B. sysctl -w net.ipv4.ip_forward=1
C. echo "net.ipv4.ip_forward=1" >> /etc/sysctl.conf
D. echo 1 > /proc/sys/net/ipv4/ip_forward
E. sysctl –p
F. echo "net.ipv6.conf.all.forwarding=l" >> /etc/sysctl.conf

Answer: BE

Explanation:
The commands that can be used together to persistently enable IPv4 forwarding in one of the Linux systems are sysctl -w net.ipv4.ip_forward=1 and sysctl –p. The first command will use sysctl to write a new value (1) to the net.ipv4.ip_forward kernel parameter, which controls whether IP forwarding is enabled or disabled for IPv4. This will enable IP forwarding immediately without rebooting. However, this change is temporary and will be lost after a reboot or a system reload. To make it permanent, we need to use the second command sysctl –p, which will load kernel parameters from /etc/sysctl.conf file. This file contains key-value pairs of kernel parameters and their values. To make sure that net.ipv4.ip_forward is set to 1 in this file, we can either edit it manually or append it using echo "net.ipv4.ip_forward=1" >> /etc/sysctl.conf.
The other options are not correct commands for persistently enabling IPv4 forwarding. The sysctl net.ipv4.ip_forward command will only display the current value of net.ipv4.ip_forward parameter, but not change it. The echo 1 > /proc/sys/net/ipv4/ip_forward command will write 1 to /proc/sys/net/ipv4/ip_forward file, which is another way to change net.ipv4.ip_forward parameter. However, this change is also temporary and will not survive a reboot or a system reload. The echo "net.ipv6.conf.all.forwarding=l" >> /etc/sysctl.conf command will append a line to /etc/sysctl.conf file that sets net.ipv6.conf.all.forwarding parameter to 1. However, this parameter controls whether IP forwarding is enabled or disabled for IPv6, not IPv4. References: sysctl(8) - Linux manual page; Configure Linux as a Router (IP Forwarding)

NEW QUESTION 197
A systems administrator is investigating an issue in which one of the servers is not booting up properly. The journalctl entries show the following:

```
Sep 16 20:30:43 server kernel: acpi PNP0A03:00: _OSC failed (AE_NOT_FOUND);
-- Subject: Unit dev-mapper-centos\x2dapp.device has failed
Sep 16 20:32:15 server systemd[1]: Dependency failed for /opt/app
-- Subject: Unit opt-app.mount has failed
-- Unit opt-app.mount has failed
Sep 16 20:32:15 server systemd[1]: Dependency failed for Local File Systems.
-- Subject: Unit local-fs.target has failed
-- Unit local-fs.target has failed.
Sep 16 20:32:15 server systemd[1]: Dependency failed for Relabel all filesystem, if necessary.
-- Subject: Unit rhel-autorelabel.service has failed
-- Unit rhel-autorelabel.service has failed.
```

Which of the following will allow the administrator to boot the Linux system to normal mode quickly?

A. Comment out the /opt/app filesystem in /etc/fstab and reboot.
B. Reformat the /opt/app filesystem and reboot.
C. Perform filesystem checks on local filesystems and reboot.
D. Trigger a filesystem relabel and reboot.

Answer: A

Explanation:
The fastest way to boot the Linux system to normal mode is to comment out the /opt/app filesystem in /etc/fstab and reboot. This will prevent the system from trying to mount the /opt/app filesystem at boot time, which causes an error because the filesystem does not exist or is corrupted. Commenting out a line in /etc/fstab can be done by adding a # symbol at the beginning of the line. Rebooting the system will apply the changes and allow the system to boot normally. Reformatting the /opt/app filesystem will not help to boot the system, as it will erase any data on the filesystem and require manual intervention to create a new filesystem. Performing filesystem checks on local filesystems will not help to boot the system, as it will not fix the missing or corrupted /opt/app filesystem. Triggering a filesystem relabel will not help to boot the system, as it will only change the security context of files and directories according to SELinux policy.
References: CompTIA Linux+ (XK0- 005) Certification Study Guide, Chapter 14: Managing Disk Storage, page 456.

NEW QUESTION 198
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

Answer: E

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). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes, page 431.

NEW QUESTION 203
A Linux administrator created the directory /project/access2all. By creating this directory, the administrator is trying to avoid the deletion or modification of files from non-owners. Which of the following will accomplish this goal?

A. chmod +t /project/access2all
B. chmod +rws /project/access2all
C. chmod 2770 /project/access2all
D. chmod ugo+rwx /project/access2all

Answer: A

Explanation:
The command that will accomplish the goal of avoiding the deletion or modification of files from non-owners is chmod +t /project/access2all. This command will set the sticky bit on the directory /project/access2all, which is a special permission that restricts file deletion or renaming to only the file owner, directory owner, or root user. This way, even if multiple users have write permission to the directory, they cannot delete or modify each other's files.
The other options are not correct commands for accomplishing the goal. The chmod +rws /project/access2all command will set both the SUID and SGID bits on the directory, which are special permissions that allow a program or a directory to run or be accessed with the permissions of its owner or group, respectively. However, this does not prevent file deletion or modification from non-owners. The chmod 2770 /project/access2all command will set only the SGID bit on the directory, which means that any new files or subdirectories created in it will inherit its group ownership. However, this does not prevent file deletion or modification from non-owners. The chmod ugo+rwx /project/access2all command will grant read, write, and execute permissions to all users (user, group, and others) on the directory, which means that anyone can delete or modify any file in it. References: chmod(1) - Linux manual page; How to Use SUID, SGID, and Sticky Bits on Linux

NEW QUESTION 208
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. systemct1 isolate multi-user.target
B. systemct1 isolate graphical.target
C. systemct1 isolate network.target
D. systemct1 isolate basic.target

Answer: B

Explanation:
The command that would ensure the server is set to runlevel 5 is systemct1 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 systemct1 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 systemct1 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 systemct1 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. References: systemd System and Service Manager; systemd.special(7) - Linux manual page

NEW QUESTION 209
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

Answer: A

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). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Users and Groups, page 463.

NEW QUESTION 214
A developer reported an incident involving the application configuration file /etc/httpd/conf/httpd.conf that is missing from the server. Which of the following identifies the RPM package that installed the configuration file?

A. rpm -qf /etc/httpd/conf/httpd.conf
B. rpm -ql /etc/httpd/conf/httpd.conf
C. rpm —query /etc/httpd/conf/httpd.conf
D. rpm -q /etc/httpd/conf/httpd.conf

Answer: A

Explanation:
The rpm -qf /etc/httpd/conf/httpd.conf command will identify the RPM package that installed the configuration file. This command will query the database of installed packages and display the name of the package that owns the specified file. The rpm -ql /etc/httpd/conf/httpd.conf command is invalid, as -ql is not a valid option for rpm. The rpm --query /etc/httpd/conf/httpd.conf command is incorrect, as --query requires a package name, not a file name. The rpm -q /etc/httpd/conf/httpd.conf command is incorrect,
as -q requires a package name, not a file name. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Packages and Software, page 560.

NEW QUESTION 218
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.

Answer: D

Explanation:
The issue is caused by the lack of a default route defined in the /etc/sysconfig/network-scripts/ifcfg-enp0s3 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. References: Configuring IP Networking with nmcli; Configuring IP Networking with ifcfg Files

NEW QUESTION 221
A Linux engineer is setting the sticky bit on a directory called devops with 755 file permission. Which of the following commands will accomplish this task?

A. chown -s 755 devops
B. chown 1755 devops
C. chmod -s 755 devops
D. chmod 1755 devops

Answer: D

Explanation:
The command that will set the sticky bit on a directory called devops with 755 file permission is chmod 1755 devops. This command will use chmod to change the mode of the directory devops to 1755, which means that the owner has read, write, and execute permissions (7), the group has read and execute permissions (5), and others have read and execute permissions (5). The first digit 1 indicates that the sticky bit is set on the directory, which is a special permission that prevents users from deleting or renaming files
in the directory that they do not own.
The other options are not correct commands for setting the sticky bit on a directory. The chown -s 755 devops command is invalid because chown is used to change the owner and group of files or directories, not their permissions. The -s option for chown is used to remove a symbolic link, not to set the sticky bit. The chown 1755 devops command is also invalid because chown does not accept numeric arguments for changing permissions. The chmod -s 755 devops command will remove the sticky bit from the directory devops, not set it. References: chmod(1) - Linux manual page; How to Use SUID, SGID, and Sticky Bits on Linux

NEW QUESTION 223
.....

Relate Links

100% Pass Your XK0-005 Exam with Exambible Prep Materials

https://www.exambible.com/XK0-005-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/