CompTIA

Exam Questions XK0-005

CompTIA Linux+ Certification Exam

NEW QUESTION 1
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 2
A systems administrator received a notification that a system is performing slowly. When running the top command, the systems administrator can see the following values:

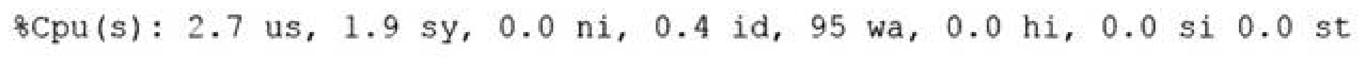

```
%Cpu(s): 2.7 us, 1.9 sy, 0.0 ni, 0.4 id, 95 wa, 0.0 hi, 0.0 si 0.0 st
```

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

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

Answer: A

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

NEW QUESTION 3
The administrator comptia is not able to perform privileged functions on a newly deployed system. Given the following command outputs:

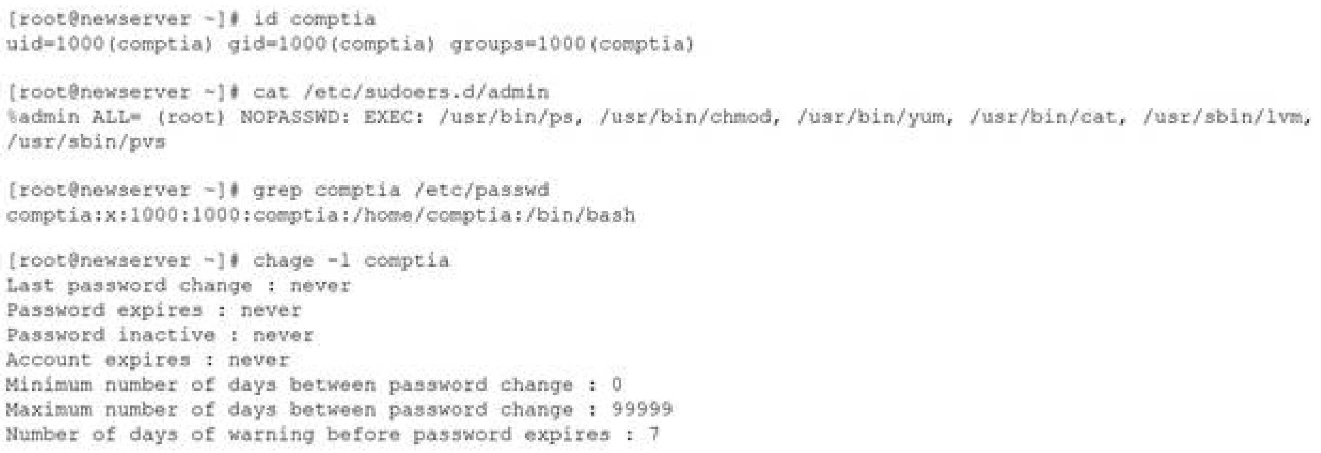

```
[root@newserver ~]# id comptia
uid=1000(comptia) gid=1000(comptia) groups=1000(comptia)

[root@newserver ~]# cat /etc/sudoers.d/admin
%admin ALL= (root) NOPASSWD: EXEC: /usr/bin/ps, /usr/bin/chmod, /usr/bin/yum, /usr/bin/cat, /usr/sbin/lvm,
/usr/sbin/pvs

[root@newserver ~]# grep comptia /etc/passwd
comptia:x:1000:1000:comptia:/home/comptia:/bin/bash

[root@newserver ~]# chage -l comptia
Last password change : never
Password expires : never
Password inactive : never
Account expires : never
Minimum number of days between password change : 0
Maximum number of days between password change : 99999
Number of days of warning before password expires : 7
```

Which of the following is the reason that the administrator is unable to perform the assigned duties?

A. The administrator needs a password reset.
B. The administrator is not a part of the correct group.
C. The administrator did not update the sudo database.
D. The administrator's credentials need to be more complex.

Answer: B

Explanation:
The reason that the administrator is unable to perform the assigned duties is because the administrator is not a part of the correct group. This is option B.
Based on the image that you sent, I can see that the user comptia has a user ID and a group ID of 1000, and belongs to only one group, which is also comptia. However, the sudoers file, which defines the permissions for users to run commands as root or other users, does not include the comptia group in any of the entries. Therefore, the user comptia cannot use sudo to perform privileged functions on the system.
The other options are incorrect because:
* A. The administrator needs a password reset.
This is not true, because the password aging information for the user comptia shows that the password was last changed on Oct 24, 2023, and it does not expire until Jan 22, 2024. There is no indication that the password is locked or expired.
* C. The administrator did not update the sudo database.
This is not necessary, because the sudo database is automatically updated whenever the sudoers file is modified. There is no separate command to update the sudo database.
* D. The administrator's credentials need to be more complex.
This is not relevant, because the complexity of the credentials does not affect the ability to use sudo. The sudoers file does not specify any password policy for the users or groups that are allowed to use sudo.

NEW QUESTION 4
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 5
In which of the following filesystems are system logs commonly stored?

A. /var
B. /tmp
C. /etc
D. /opt

Answer: A

Explanation:
The filesystem that system logs are commonly stored in is /var. The /var filesystem is a directory that contains variable data files on Linux systems. Variable data files are files that are expected to grow in size over time, such as logs, caches, spools, and temporary files. The /var filesystem is separate from the / filesystem, which contains the essential system files, to prevent the / filesystem from being filled up by the variable data files. The system logs are files that record the events and activities of the system and its components, such as the kernel, the services, the applications, and the users. The system logs are useful for monitoring, troubleshooting, and auditing the system. The system logs are commonly stored in the /var/log directory, which is a subdirectory of the /var filesystem. The /var/log directory contains various log files, such as syslog, messages, dmesg, auth.log, and kern.log. The filesystem that system logs are commonly stored in is /var. This is the correct answer to the question. The other options are incorrect because they are not the filesystems that system logs are commonly stored in (/tmp, /etc, or /opt). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Managing Logging and Monitoring, page 487.

NEW QUESTION 6
An administrator recently updated the BIND software package and would like to review the default configuration that shipped with this version. Which of the following files should the administrator review?

A. /etc/named.conf.rpmnew
B. /etc/named.conf.rpmsave
C. /etc/named.conf
D. /etc/bind/bind.conf

Answer: A

Explanation:
After installing a new version of a package that includes a configuration file that already exists on the system, such as /etc/httpd/conf/httpd.conf, RPM will create a new file with the .rpmnew extension instead of overwriting the existing file. This allows the administrator to review the default configuration that shipped with this version and compare it with the current configuration before deciding whether to merge or replace the files. The /etc/named.conf.rpmsave file is created by RPM when a package is uninstalled and it contains a configuration file that was modified by the administrator. This allows the administrator to restore the configuration file if needed. The /etc/named.conf file is the main configuration file for the BIND name server, not the httpd web server. The /etc/bind/bind.conf file does not exist by default in Linux systems. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Packages and Software, page 561.

NEW QUESTION 7
A non-privileged user is attempting to use commands that require elevated account permissions, but the commands are not successful. Which of the following most likely needs to be updated?

A. /etc/passwd
B. /etc/shadow
C. /etc/sudoers
D. /etc/bashrc

Answer: C

Explanation:
The /etc/sudoers file is used to configure the sudo command, which allows non-privileged users to execute commands that require elevated account permissions1. The file contains a list of users and groups that are allowed to use sudo, and the commands they can run with it. The file also defines the security policy for sudo, such as whether a password is required, how long the sudo session lasts, and what environment variables are preserved or reset.
The /etc/passwd file is used to store information about the user accounts on the system, such as their username, user ID, home directory, and login shell. The /etc/shadow file is used to store the encrypted passwords for the user accounts, along with other information such as password expiration and aging. These files are not directly related to the sudo command, and updating them will not grant a user elevated account permissions.
The /etc/bashrc file is used to set up the environment for the bash shell, such as aliases, functions, variables, and options. This file is executed whenever a new

bash shell is started, and it affects all users on the system. However, this file does not control the sudo command or its configuration, and updating it will not allow a user to use commands that require elevated account permissions.

NEW QUESTION 8
A Linux administrator is troubleshooting a systemd mount unit file that is not working correctly. The file contains:
[root@system] # cat mydocs.mount [Unit]
Description=Mount point for My Documents drive [Mount]
What=/dev/drv/disk/by-uuid/94afc9b2-ac34-ccff-88ae-297ab3c7ff34 Where=/home/user1/My Documents
Options=defaults Type=xfs
[Install]
WantedBy=multi-user.target
The administrator verifies the drive UUID correct, and user1 confirms the drive should be
mounted as My Documents in the home directory. Which of the following can the administrator
do to fix the issues with mounting the drive? (Select two).

A. Rename the mount file to home-user1-My\x20Documents.mount.
B. Rename the mount file to home-user1-my-documents.mount.
C. Change the What entry to /dev/drv/disk/by-uuid/94afc9b2\-ac34\-ccff\-88ae\- 297ab3c7ff34.
D. Change the Where entry to Where=/home/user1/my\ documents.
E. Change the Where entry to Where=/home/user1/My\x20Documents.
F. Add quotes to the What and Where entries, such as What="/dev/drv/disk/by- uuid/94afc9b2-ac34-ccff-88ae-297ab3c7ff34" and Where="/home/user1/My Documents".

Answer: AE

Explanation:
The mount unit file name and the Where entry must be escaped to handle spaces in the path.ReferencesThe mount unit file name must be named after the mount point directory, with spaces replaced by \x20. See How to escape spaces in systemd unit files? and systemd.mount.The Where entry must use \x20 to escape spaces in the path. See systemd.mount and The workaround is to use /usr/bin/env followed by the path in quotes..

NEW QUESTION 9
During a security scan, the password of an SSH key file appeared to be too weak and was cracked. Which of the following commands would allow a user to choose a stronger password and set it on the existing SSH key file?

A. passwd
B. ssh
C. ssh-keygen
D. pwgen

Answer: C

Explanation:
The command that would allow a user to choose a stronger password and set it on the existing SSH key file is ssh-keygen -p -f <keyfile>. This command uses the ssh-keygen tool, which is used to generate, manage, and convert authentication keys for SSH. The -p option stands for passphrase, and it allows the user to change or remove the passphrase of an existing private key file. The -f option specifies the filename of the key file. The command will prompt the user for the old passphrase, and then for the new passphrase twice.
The other options are not correct commands for changing the password of an SSH key file. The passwd command is used to change the password of a user account on a Linux system, not an SSH key file. The ssh command is used to log in to a remote system using SSH, not to change the password of an SSH key file. The pwgen command is used to generate random passwords, not to change the password of an SSH key file.
References: ssh-keygen(1) - Linux manual page; How To: Change Passphrase for SSH Private Key - Unix Tutorial

NEW QUESTION 10
A systems administrator wants to permit access temporarily to an application running on port 1234/TCP on a Linux server. Which of the following commands will permit this traffic?

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

Answer: C

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

NEW QUESTION 10
A Linux administrator was asked to run a container with the httpd server inside. This container should be exposed at port 443 of a Linux host machine while it internally listens on port 8443. Which of the following commands will accomplish this task?

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

C. podman run –d -e 443:8443 httpd
D. podman exec -p 8443:443 httpd

Answer: A

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

NEW QUESTION 14
A Linux administrator is troubleshooting a memory-related issue. Based on the output of the commands:

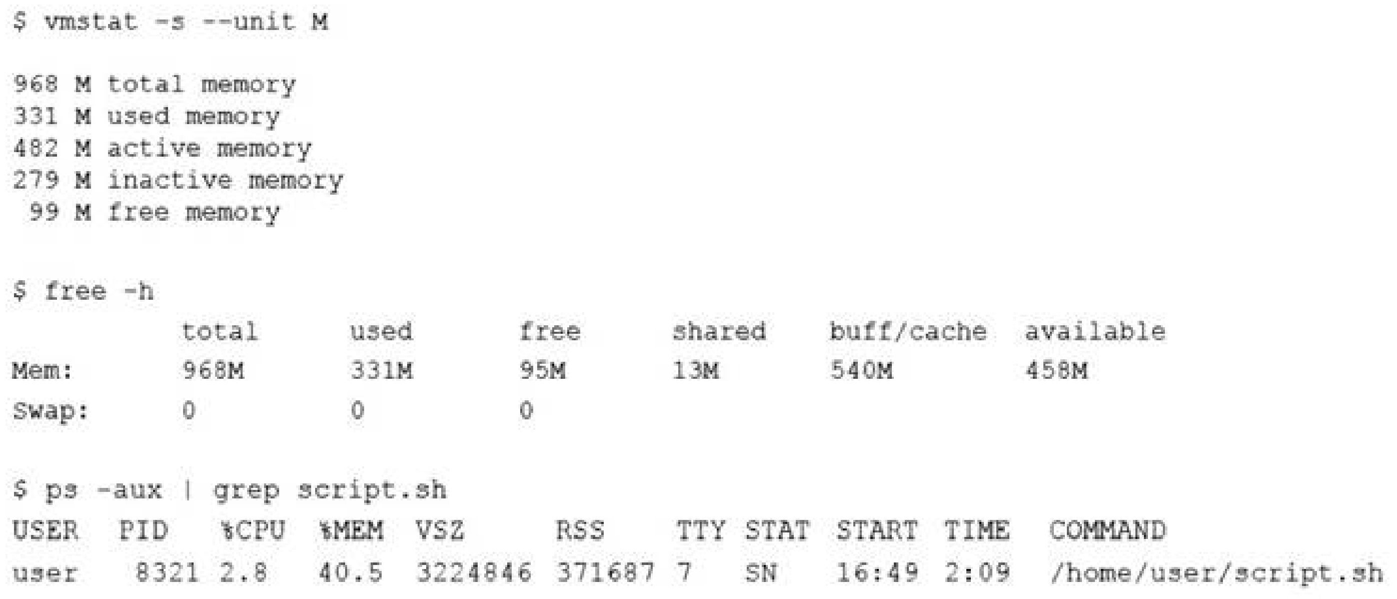

```
$ vmstat -s --unit M

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

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

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

Which of the following commands would address the issue?

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

Answer: B

Explanation:
The command that would address the memory-related issue is kill -9 8321. This command will send a SIGKILL signal to the process with the PID 8321, which is the mysqld process that is using 99.7% of the available memory according to the top output. The SIGKILL signal will terminate the process immediately and free up the memory it was using. However, this command should be used with caution as it may cause data loss or corruption if the process was performing some critical operations.
The other options are not correct commands for addressing the memory-related issue. The top -p 8321 command will only display information about the process with the PID 8321, but will not kill it or reduce its memory usage. The renice -10 8321 command will change the priority (niceness) of the process with the PID 8321 to -10, which means it will have a higher scheduling priority, but this will not affect its memory consumption. The free 8321 command is invalid because free does not take a PID as an argument; free only displays information about the total, used, and free memory in the system. References: How to troubleshoot Linux server memory issues; kill(1) - Linux manual page

NEW QUESTION 17
A systems technician is working on deploying several microservices to various RPM-based systems, some of which could run up to two hours. Which of the following commands will allow the technician to execute those services and continue deploying other microservices within the same terminal section?

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

Answer: D

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

NEW QUESTION 18
A systems administrator is tasked with setting up key-based SSH authentication. In which of the following locations should the administrator place the public keys for the server?

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

Answer: C

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

NEW QUESTION 19
To harden one of the servers, an administrator needs to remove the possibility of remote administrative login via the SSH service. Which of the following should the administrator do?

A. Add the line DenyUsers root to the /etc/hosts.deny file.
B. Set PermitRootLogin to no in the /etc/ssh/sshd_config file.
C. Add the line account required pam_nologi
D. so to the /etc/pam.d/sshd file.
E. Set PubKeyAuthentication to no in the /etc/ssh/ssh_config file.

Answer: B

Explanation:
The administrator should set PermitRootLogin to no in the /etc/ssh/sshd_config file to remove the possibility of remote administrative login via the SSH service. The PermitRootLogin directive controls whether the root user can log in using SSH. Setting it to no will deny any remote login attempts by the root user. This will harden the server and prevent unauthorized access. The administrator should also restart the sshd service after making the change. The other options are incorrect because they either do not affect the SSH service (/etc/hosts.deny or /etc/pam.d/sshd) or do not prevent remote administrative login (PubKeyAuthentication). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 413.

NEW QUESTION 23
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 28
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 32
A systems administrator requires that all files that are created by the user named web have read-only permissions by the owner. Which of the following commands will satisfy this requirement?

A. chown web:web /home/web
B. chmod -R 400 /home/web
C. echo "umask 377" >> /home/web/.bashrc
D. setfacl read /home/web

Answer: C

Explanation:
The command that will satisfy the requirement of having all files that are created by the user named web have read-only permissions by the owner is echo "umask 377" >> /home/web/.bashrc. This command will append the umask 377 command to the end of the .bashrc file in the web user's home directory. The .bashrc file is a shell script that is executed whenever a new interactive shell session is started by the user. The umask command sets the file mode creation mask, which determines the default permissions for newly created files or directories by subtracting from the maximum permissions (666 for files and 777 for directories). The umask 377 command means that the user does not want to give any permissions to the group or others (3 = 000 in binary), and only wants to give read permission to the owner (7 - 3 = 4 = 100 in binary). Therefore, any new file created by the web user will have read-only permission by the owner (400) and no permission for anyone else. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing Users and Groups; Umask Command in Linux | Linuxize

NEW QUESTION 33
Users are reporting that writes on a system configured with SSD drives have been taking longer than expected, but reads do not seem to be affected. A Linux systems administrator is investigating this issue and working on a solution. Which of the following should the administrator do to help solve the issue?

A. Run the corresponding command to trim the SSD drives.
B. Use fsck on the filesystem hosted on the SSD drives.
C. Migrate to high-density SSD drives for increased performance.
D. Reduce the amount of files on the SSD drives.

Answer: A

Explanation:
TRIM is a feature that allows the operating system to inform the SSD which blocks of data are no longer in use and can be wiped internally. This helps to maintain the SSD's performance and endurance by preventing unnecessary write operations and reducing write amplification12. Running the corresponding command to trim the SSD drives, such as fstrim or blkdiscard on Linux, can help to solve the issue of slow writes by freeing up space and optimizing the SSD's internal garbage collection34.
References: 1: What is SSD TRIM, why is it useful, and how to check whether it is turned on 2: How to Trim SSD in Windows 10 3: How to run fsck on an external drive with OS X? 4: How to Use the fsck Command on Linux

NEW QUESTION 38
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 43
A DevOps engineer wants to allow the same Kubernetes container configurations to be deployed in development, testing, and production environments. A key requirement is that the containers should be configured so that developers do not have to statically configure custom, environment-specific locations. Which of the following should the engineer use to meet this requirement?

A. Custom scheduler
B. Node affinity
C. Overlay network
D. Ambassador container

Answer: D

Explanation:
To allow the same Kubernetes container configurations to be deployed in different environments without statically configuring custom locations, the engineer can use an ambassador container (D). An ambassador container is a proxy container that handles communication between containers and external services. It can dynamically configure locations based on environment variables or other methods. The other options are not related to this requirement. References:
? [CompTIA Linux+ Study Guide], Chapter 11: Working with Containers, Section: Using Ambassador Containers
? [How to Use Ambassador Containers]

NEW QUESTION 45
Which of the following can be used as a secure way to access a remote termi-nal?

A. TFTP
B. SSH
C. SCP
D. SFTP

Answer: B

Explanation:
SSH, or Secure Shell, is a protocol that allows you to access a remote terminal or virtual machine securely over an encrypted connection. You can use SSH to run

commands, transfer files, or tunnel network traffic on a remote system. To use SSH, you need an SSH client program on your local system and an SSH server program on the remote system. You also need to authenticate yourself using a username and password or a public/private key pair. SSH is widely used by system administrators, developers, and engineers to remotely manage Linux servers and other devices.
The other options are not correct answers. TFTP, or Trivial File Transfer Protocol, is a simple protocol that allows you to transfer files between systems, but it does not provide any security or encryption features. SCP, or Secure Copy Protocol, is a protocol that uses SSH to securely copy files between systems, but it does not provide a remote terminal access. FTP, or File Transfer Protocol, is another protocol that allows you to transfer files between systems, but it also does not provide any security or encryption features.

NEW QUESTION 48
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 53
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:

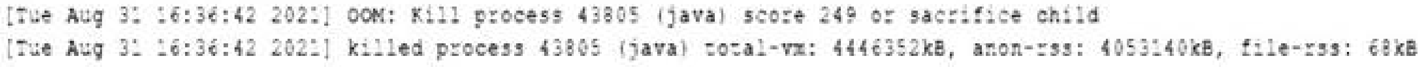

```
[Tue Aug 31 16:36:42 2021] OOM: Kill process 43805 (java) score 249 or sacrifice child
[Tue Aug 31 16:36:42 2021] killed process 43805 (java) total-vm: 4446352kB, anon-rss: 4053140kB, file-rss: 68kB
```

Output 2:

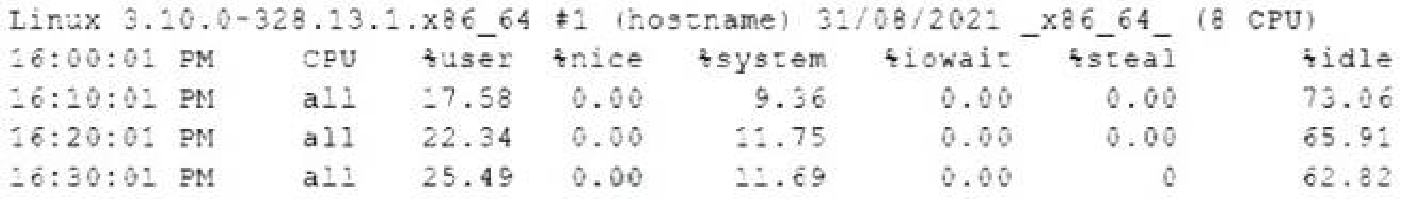

```
Linux 3.10.0-328.13.1.x86_64 #1 (hostname) 31/08/2021 _x86_64_ (8 CPU)
16:00:01 PM     CPU    %user   %nice   %system   %iowait   %steal      %idle
16:10:01 PM     all    17.58    0.00      9.36      0.00     0.00      73.06
16:20:01 PM     all    22.34    0.00     11.75      0.00     0.00      65.91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

Output 3:

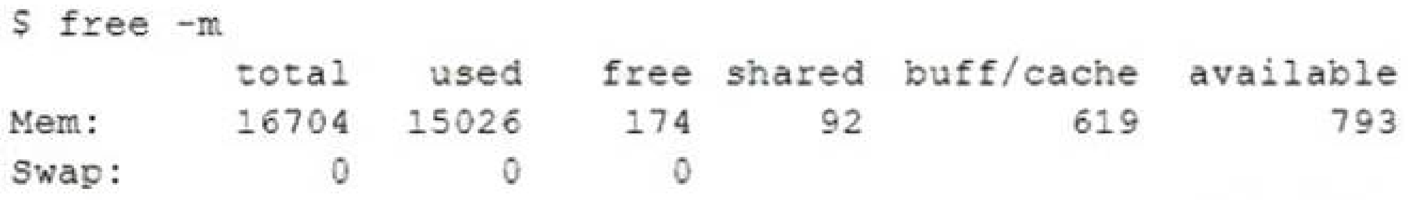

```
$ free -m
          total    used    free  shared  buff/cache  available
Mem:      16704   15026     174      92         619        793
Swap:         0       0       0
```

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 54
A Linux engineer set up two local DNS servers (10.10.10.10 and 10.10.10.20) and was testing email connectivity to the local mail server using the mail command on a local machine when the following error appeared:

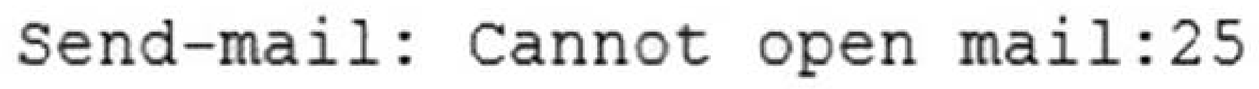

```
Send-mail: Cannot open mail:25
```

The local machine DNS settings are:

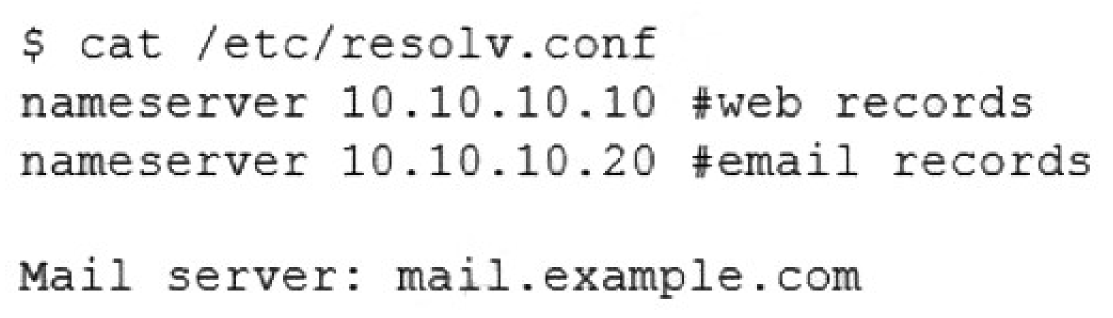

```
$ cat /etc/resolv.conf
nameserver 10.10.10.10 #web records
nameserver 10.10.10.20 #email records

Mail server: mail.example.com
```

Which of the following commands could the engineer use to query the DNS server to get mail server information?

A. dig @example.com 10.10.10.20 a
B. dig @10.10.10.20 example.com mx
C. dig @example.com 10.10.10.20 ptr
D. dig @10.10.10.20 example.com ns

Answer: B

Explanation:
The command dig @10.10.10.20 example.com mx will query the DNS server to get mail server information. The dig command is a tool for querying DNS servers and displaying the results. The @ option specifies the DNS server to query, in this case 10.10.10.20. The mx option specifies the type of record to query, in this case mail exchange (MX) records, which identify the mail servers for a domain. The domain name to query is example.com. This command will show the MX records for example.com from the DNS server 10.10.10.20. This is the correct command to use to accomplish the task. The other options are incorrect because they either use the wrong syntax (@example.com 10.10.10.20 instead of @10.10.10.20 example.com), the wrong type of record (a or ptr instead of mx), or the wrong domain name (example.com ns instead of example.com mx). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 415.

NEW QUESTION 55
Which of the following is the best tool for dynamic tuning of kernel parameters?

A. tuned
B. tune2fs
C. tuned-adm
D. turbostat

Answer: A

Explanation:
The tuned application is the best tool for dynamic tuning of kernel parameters, as it monitors the system and optimizes the performance under different workloads. It provides a number of predefined profiles for typical use cases, such as power saving, low latency, high throughput, virtual machine performance, and so on. It also allows users to create, modify, and delete profiles, and to switch between them on the fly. The tuned application uses the sysctl command and the configuration files in the /etc/sysctl.d/ directory to adjust the kernel parameters at runtime.
References
? Chapter 2. Getting started with TuneD - Red Hat Customer Portal, paragraph 1
? Kernel tuning with sysctl - Linux.com, paragraph 1

NEW QUESTION 59
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 61
A Linux administrator copied a Git repository locally, created a feature branch, and committed some changes to the feature branch. Which of the following Git actions should the Linux administrator use to publish the changes to the main branch of the remote repository?

A. rebase
B. tag
C. commit
D. push

Answer: D

Explanation:
The push action is used to publish the changes made in a local branch to a remote branch of a Git repository. This action will update the remote branch with the commits made in the local branch and synchronize the two branches. The rebase action is used to reapply commits from one branch onto another branch, creating a linear history of commits. This action does not publish any changes to a remote repository. The tag action is used to create an annotated reference to a specific commit in a Git repository. This action does not publish any changes to a remote repository. The commit action is used to record changes made in the local repository and create a new snapshot of the project state. This action does not publish any changes to a remote repository. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 20: Writing and Executing Bash Shell Scripts, page 579.

NEW QUESTION 63
An administrator needs to make some changes in the IaC declaration templates. Which of the following commands would maintain version control?

A. git clone https://github.com/comptia/linux+-.git git push origin
B. git clone https://qithub.com/comptia/linux+-.git git fetch New-Branch
C. git clone https://github.com/comptia/linux+-.git git status
D. git clone https://github.com/comptia/linuxt+-.git git checkout -b <new-branch>

Answer: D

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

NEW QUESTION 66
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 68
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 73
A systems administrator is tasked with installing GRUB on the legacy MBR of the SATA hard drive. Which of the following commands will help the administrator accomplish this task?

A. grub-install /dev/hda
B. grub-install /dev/sda
C. grub-install /dev/sr0
D. grub-install /dev/hd0,0

Answer: B

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

NEW QUESTION 78
An administrator would like to securely connect to a server and forward port 8080 on a local machine to port 80 on the server. Which of the following commands should the administrator use to satisfy both requirements?

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

Answer: A

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

NEW QUESTION 81
A systems administrator is compiling a report containing information about processes that are listening on the network ports of a Linux server. Which of the following commands will allow the administrator to obtain the needed information?

A. ss -pint
B. tcpdump -nL
C. netstat -pn
D. lsof -lt

Answer: A

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

NEW QUESTION 84
Due to performance issues on a server, a Linux administrator needs to termi-nate an unresponsive process. Which of the following commands should the administrator use to terminate the process immediately without waiting for a graceful shutdown?

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

Answer: A

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

? [How to Kill Processes in Linux]

NEW QUESTION 89

A Linux systems administrator is setting up a new web server and getting 404 - NOT FOUND errors while trying to access the web server pages from the browser. While working on the diagnosis of this issue, the Linux systems administrator executes the following commands:

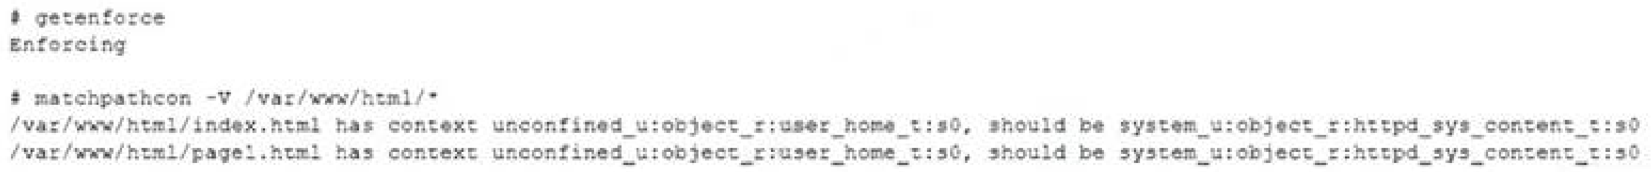

```
# getenforce
Enforcing

# matchpathcon -V /var/www/html/*
/var/www/html/index.html has context unconfined_u:object_r:user_home_t:s0, should be system_u:object_r:httpd_sys_content_t:s0
/var/www/html/page1.html has context unconfined_u:object_r:user_home_t:s0, should be system_u:object_r:httpd_sys_content_t:s0
```

Which of the following commands will BEST resolve this issue?

A. sed -i 's/SELINUX=enforcing/SELINUX=disabled/' /etc/selinux/config
B. restorecon -R -v /var/www/html
C. setenforce 0
D. setsebool -P httpd_can_network_connect_db on

Answer: B

Explanation:

The command restorecon -R -v /var/www/html will best resolve the issue. The issue is caused by the incorrect SELinux context of the web server files under th /var/www/html directory. The output of ls -Z /var/www/html shows that the files have the type user_home_t, which is not allowed for web content. The command restorecon restores the default SELinux context of files based on the policy rules. The options -R and -v are used to apply the command recursively and verbosely. This command will change the type
of the files to httpd_sys_content_t, which is the correct type for web content. This will allow the web server to access the files and serve the pages to the browser. The other options are incorrect because they either disable SELinux entirely (sed -i 's/SELINUX=enforcing/SELINUX=disabled/' /etc/selinux/config or setenforce 0), which is not a good security practice, or enable an unnecessary boolean (setsebool -P httpd_can_network_connect_db on), which is not related to the issue. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 535.

NEW QUESTION 91

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 95

A user created the following script file:
# ! /bin/bash
# FILENAME: /home/user/ script . sh echo "hello world"
exit 1
However, when the user tried to run the script file using the command "script . sh, an error returned indicating permission was denied. Which of the follow-ing should the user execute in order for the script to run properly?

A. chmod u+x /home/user/script . sh
B. chmod 600 /home/user/script . sh
C. chmod /home/user/script . sh
D. chmod 0+r /horne/user/scrip
E. sh

Answer: A

Explanation:

To run a script file, the user needs to have execute permission on the file. The command chmod u+x /home/user/script.sh (A) will grant execute permission to the owner of the file, which is the user who created it. The other commands will not give execute permission to the user, and therefore will not allow the script to run properly. References:
? [CompTIA Linux+ Study Guide], Chapter 3: Working with Files, Section: Changing File Permissions
? [How to Make a Bash Script Executable]

NEW QUESTION 99

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

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

Answer: A

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

NEW QUESTION 101
An administrator deployed a Linux server that is running a web application on port 6379/tcp.
SELinux is in enforcing mode based on organization policies. The port is open on the firewall.
Users who are trying to connect to a local instance of the web application receive Error 13, Permission denied.
The administrator ran some commands that resulted in the following output:

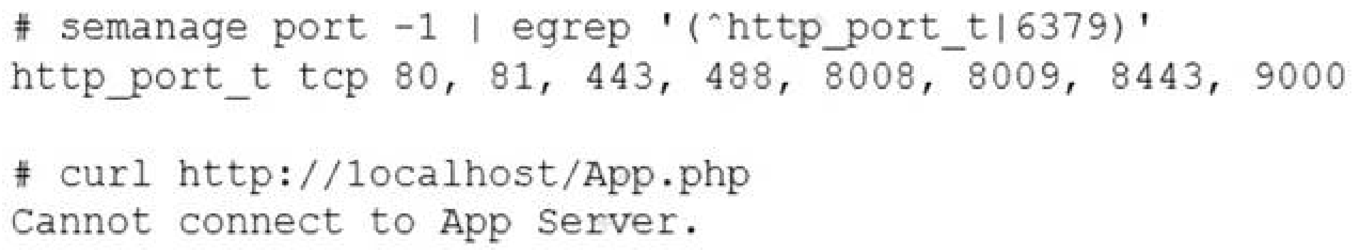

```
# semanage port -l | egrep '(^http_port_t|6379)'
http_port_t tcp 80, 81, 443, 488, 8008, 8009, 8443, 9000

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

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

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

Answer: B

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

NEW QUESTION 102
A Linux administrator has set up a new DNS forwarder and is configuring all internal servers to use the new forwarder to look up external DNS requests. The administrator needs to modify the firewall on the server for the DNS forwarder to allow the internal servers to communicate to it and make the changes persistent between server reboots. Which of the following commands should be run on the DNS forwarder server to accomplish this task?

A. ufw allow out dns
B. systemct1 reload firewalld
C. iptables -A OUTPUT -p udp -ra udp -dport 53 -j ACCEPT
D. flrewall-cmd --zone-public --add-port-53/udp --permanent

Answer: D

Explanation:
The command that should be run on the DNS forwarder server to
accomplish the task is firewall-cmd --zone=public --add-port=53/udp --permanent.
The firewall-cmd command is a tool for managing firewalld, which is a firewall service that provides dynamic and persistent network security on Linux systems. The firewalld uses zones and services to define the rules and policies for the network traffic. The zones are logical groups of network interfaces and sources that have the same level of trust and security. The services are predefined sets of ports and protocols that are associated with certain applications or functions. The --zone=public option specifies the zone name that the rule applies to. The public zone is the default zone that represents the untrusted network, such as the internet. The --add-port=53/udp option adds a port and protocol to the zone. The 53 is the port number that is used by the DNS service. The udp is the protocol that is used by the DNS service. The --permanent option makes the change persistent across reboots. The command firewall-cmd --zone=public --add-port=53/udp --permanent will modify the firewall on the server for the DNS forwarder to allow the internal servers to communicate to it and make the changes persistent between server reboots. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not modify the firewall on the server for the DNS forwarder (ufw allow out dns or systemct1 reload firewalld) or do not use the correct syntax for the command (iptables -A OUTPUT -p udp -ra udp -dport 53 -j ACCEPT instead of iptables -A OUTPUT - p udp -ra udp --dport 53 -j ACCEPT). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 392.

NEW QUESTION 105
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 106
A Linux administrator is troubleshooting an issue in which an application service failed to start on a Linux server. The administrator runs a few commands and gets the following outputs:

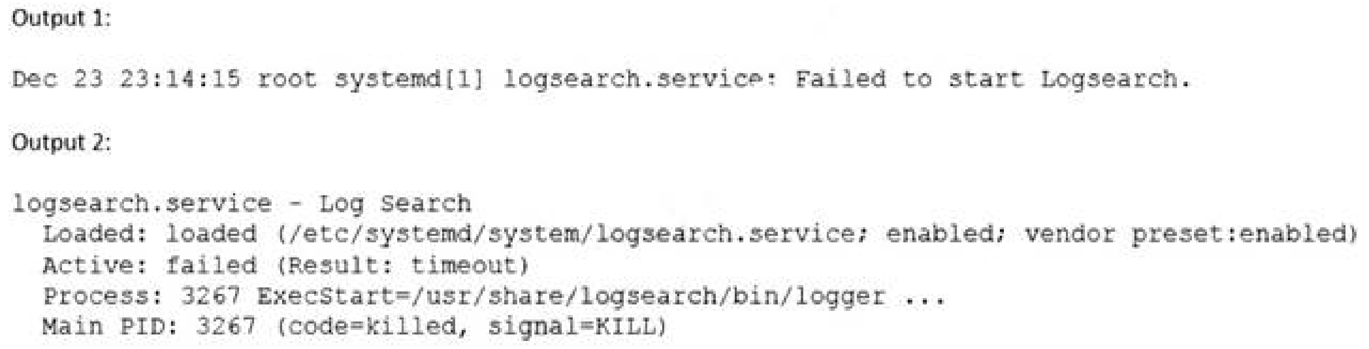

Output 1:

```
Dec 23 23:14:15 root systemd[1] logsearch.service: Failed to start Logsearch.
```

Output 2:

```
logsearch.service - Log Search
  Loaded: loaded (/etc/systemd/system/logsearch.service; enabled; vendor preset:enabled)
  Active: failed (Result: timeout)
  Process: 3267 ExecStart=/usr/share/logsearch/bin/logger ...
  Main PID: 3267 (code=killed, signal=KILL)
```

Based on the above outputs, which of the following is the MOST likely action the administrator should take to resolve this issue?

A. Enable the logsearch.service and restart the service.
B. Increase the TimeoutStartUSec configuration for the logsearch.sevice.
C. Update the OnCalendar configuration to schedule the start of the logsearch.service.
D. Update the KillSignal configuration for the logsearch.service to use TERM.

Answer: B

Explanation:
The administrator should increase the TimeoutStartUSec configuration for the logsearch.service to resolve the issue. The output of systemct1 status logsearch.service shows that the service failed to start due to a timeout. The output of cat /etc/systemd/system/logsearch.service shows that the service has a TimeoutStartUSec configuration of 10 seconds, which might be too short for the service to start. The administrator should increase this value to a higher number, such as 30 seconds or 1 minute, and then restart the service. The other options are incorrect because they are not related to the issue. The service is already enabled, as shown by the output of systemct1 is-enabled logsearch.service. The service does not use an OnCalendar configuration, as it is not a timer unit. The service does not use a KillSignal configuration, as it is not being killed by a signal. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, pages 434-435.

NEW QUESTION 109
A Linux administrator reviews a set of log output files and needs to identify files that contain any occurrence of the word denied. All log files containing entries in uppercase or lowercase letters should be included in the list. Which of the following commands should the administrator use to accomplish this task?

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

Answer: D

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

NEW QUESTION 112
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 114
A Linux administrator recently downloaded a software package that is currently in a compressed file. Which of the following commands will extract the files?

A. unzip -v
B. bzip2 -z
C. gzip
D. funzip

Answer: C

Explanation:
The command gzip can extract files that are compressed with the gzip format, which has the extension .gz. This is the correct command to use for the software package. The other options are incorrect because they either compress files (bzip2 -z), unzip files that are compressed with the zip format (unzip -v or funzip), or have the wrong options (-v or -z instead of -d). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, page 353.

NEW QUESTION 119
A Linux system is having issues. Given the following outputs:

```
# dig @192.168.2.2 mycomptiahost
; << >> DiG 9.9.4-RedHat-9.9.4-74.el7_6.1 << >> @192.168.2.2 mycomptiahost
; (1 server found)
;; global options: +cmd
;; connection timed out; no servers could be reached
# nc -v 192.168.2.2 53
Ncat: Version 7.70 ( https://nmap.org/ncat ) Ncat: Connection timed out.
# ping 192.168.2.2
PING 192.168.2.2 (192.168.2.2) 56(84) bytes of data.
64 bytes from 192.168.2.2: icmp_seq=1 ttl=117 time=4.94 ms 64 bytes from 192.168.2.2: icmp_seq=2 ttl=117 time=10.5 ms
```

Which of the following best describes this issue?

A. The DNS host is down.
B. The name mycomptiahost does not exist in the DNS.
C. The Linux engineer is using the wrong DNS port.
D. The DNS service is currently not available or the corresponding port is blocked.

Answer: D

Explanation:
The ping command shows that the Linux system can reach the DNS server at 192.168.2.2, so the DNS host is not down. The dig and nc commands show that the Linux system cannot connect to the DNS server on port 53, which is the standard port for DNS queries. This means that either the DNS service is not running on the DNS server, or there is a firewall or network device blocking the port 53 traffic. Therefore, the DNS service is currently not available or the corresponding port is blocked.References1: How To Troubleshoot DNS Client Issues in Linux - RootUsers2: 6 Best Tools to Troubleshoot DNS Issues in Linux - Tecmint3: How To Troubleshoot DNS in Linux - OrcaCore4: Fixing DNS Issues in Ubuntu 20.04 | DeviceTests

NEW QUESTION 120
A systems administrator is deploying three identical, cloud-based servers. The administrator is using the following code to complete the task:

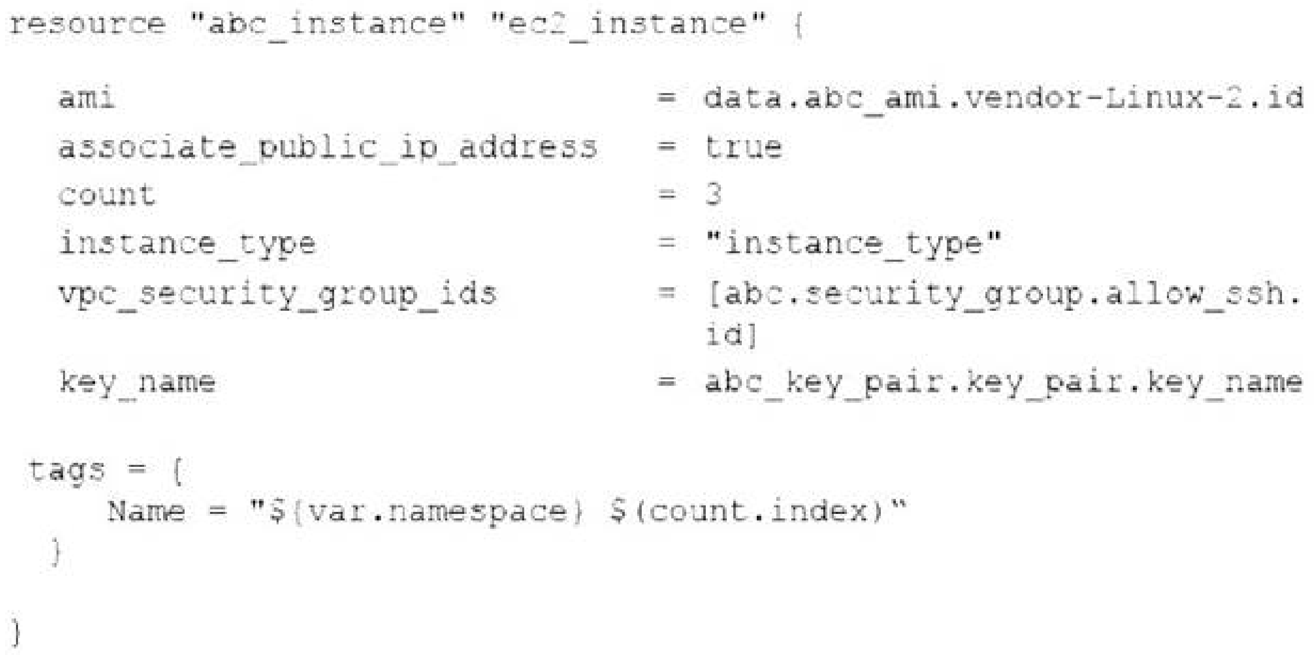

```
resource "abc_instance" "ec2_instance" {

   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 123
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 124
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:

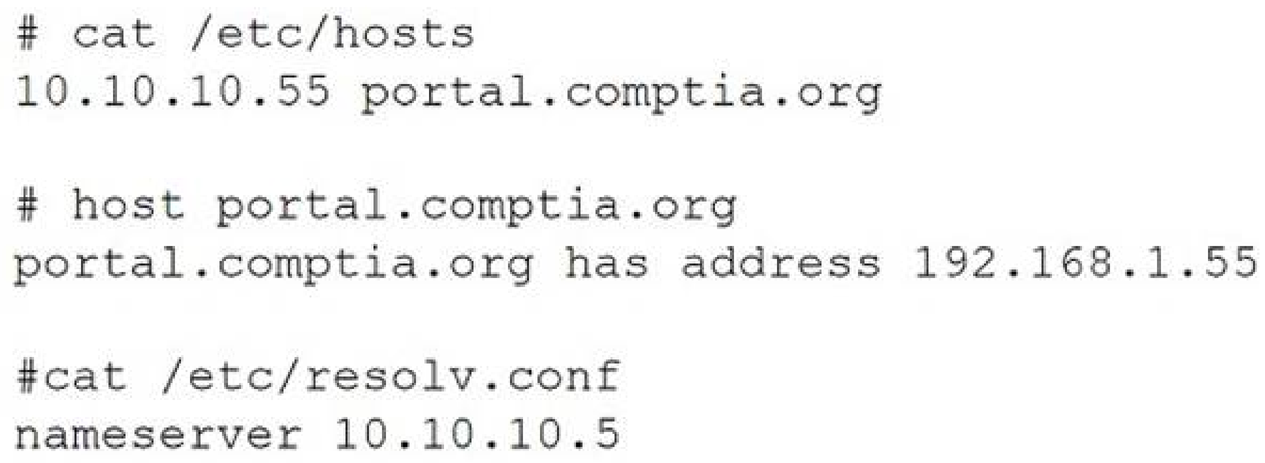

```
# cat /etc/hosts
10.10.10.55 portal.comptia.org

# 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 126
A developer wants to ensure that all files and folders created inside a shared folder named
/GroupOODEV inherit the group name of the parent folder. Which of the following commands will help achieve this goal?

A. chmod g+X / GroupOODEV/
B. chmod g+W / GroupOODEV/
C. chmod g+r / GroupOODEV/
D. chmod g+s / GroupOODEV/

Answer: D

Explanation:
The chmod command is used to change the permissions of files and directories on Linux systems. The g+s option sets the setgid bit on a directory, which means that all files and folders created inside that directory will inherit the group name of the parent directory. This command can help the developer ensure that all files and folders created inside the /GroupOODEV directory have the same group name as /GroupOODEV. References: [How to Use chmod Command in Linux with Examples]

NEW QUESTION 128
Which of the following files holds the system configuration for journal when running systemd?

A. /etc/systemd/journald.conf
B. /etc/systemd/systemd-journalctl.conf
C. /usr/lib/systemd/journalctl.conf
D. /etc/systemd/systemd-journald.conf

Answer: A

Explanation:
The file that holds the system configuration for journal when running systemd is /etc/systemd/journald.conf. This file contains various settings that control the behavior of the journald daemon, which is responsible for collecting and storing log messages from various sources. The journald.conf file can be edited to change the default values of these settings, such as the storage location, size limits, compression, and forwarding options of the journal files. The file also supports a drop-in directory /etc/systemd/journald.conf.d/ where additional configuration files can be placed to override or extend the main file. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks; journald.conf(5) - Linux manual page

NEW QUESTION 133
A systems administrator is tasked with changing the default shell of a system account in order to disable iterative logins. Which of the following is the best option for the administrator to use as the new shell?

A. /sbin/nologin
B. /bin/ sh
C. /sbin/ setenforce
D. /bin/bash

Answer: A

Explanation:
The /sbin/nologin shell is a special shell that prevents the user from logging into an interactive session. It is commonly used for system accounts that are not meant to be accessed by users, such as daemon or service accounts. When a user tries to log in with this shell, they will see a message like “This account is currently not available” and the login will fail.
References:
? The /sbin/nologin shell is listed as one of the valid shells in the /etc/shells file1.
? The CompTIA Linux+ Certification Exam Objectives mention that the candidate should be able to “configure and manage system accounts and groups, including password aging and restricted shells” as part of the Hardware and System Configuration domain2.
? The usermod command can be used to change the user’s login shell with the -s or --shell option3. For example, to change the shell of a user named daemon to /sbin/nologin, the command would be: sudo usermod -s /sbin/nologin daemon

NEW QUESTION 137
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 141
Developers have requested implementation of a persistent, static route on the application server. Packets sent over the interface eth0 to 10.0.213.5/32 should be routed via 10.0.5.1. Which of the following commands should the administrator run to achieve this goal?

A. route -i etho -p add 10.0.213.5 10.0.5.1
B. route modify eth0 +ipv4.routes "10.0.213.5/32 10.0.5.1"
C. echo "10.0.213.5 10.0.5.1 eth0" > /proc/net/route
D. ip route add 10.0.213.5/32 via 10.0.5.1 dev eth0

Answer: D

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

NEW QUESTION 143
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:

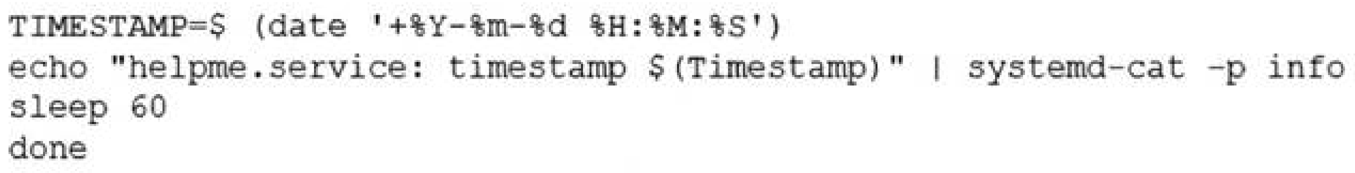

```
TIMESTAMP=$ (date '+%Y-%m-%d %H:%M:%S')
echo "helpme.service: timestamp $(Timestamp)" | systemd-cat -p info
sleep 60
done
```

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 145
An administrator would like to mirror the website files on the primary web server, www1, to the backup web server, www2. Which of the following commands should the administrator use to most efficiently accomplish this task?

A. [wwwl] rsync —a —e ssh /var/www/html/ user1@www2 : /var/www/html
B. [wwwl] scp —r /var/www/html user1@www2 : / var/www/html
C. [www2] cd /var/www/html; wget —m http: //wwwl/
D. [wwwl] cd /var/www/html && tar cvf —

Answer: A

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

NEW QUESTION 149
An administrator started a long-running process in the foreground that needs to continue without interruption. Which of the following keystrokes should the administrator use to continue running the process in the background?

A. <Ctrl+z> bg
B. <Ctrl+d> bg
C. <Ctrl+b> jobs -1
D. <Ctrl+h> bg &

Answer: A

Explanation:
A long-running process is a program that takes a long time to complete or runs indefinitely on a Linux system. A foreground process is a process that runs in the current terminal and receives input from the keyboard and output to the screen. A background process is a process that runs in the background and does not interact with the terminal. A background process can continue running even if the terminal is closed or disconnected.
To start a long-running process in the background, the user can append an ampersand (&)
to the command, such as someapp &. This will run someapp in the background and return control to the terminal immediately.
To move a long-running process from the foreground to the background, the user can use two keystrokes: Ctrl+Z and bg. The Ctrl+Z keystroke will suspend (pause) the foreground process and return control to the terminal. The bg keystroke will resume (continue) the suspended process in the background and detach it from the terminal. The statement B is correct.
The statements A, C, and D are incorrect because they do not perform the desired task. The bg keystroke alone will not work unless there is a suspended process to resume. The Ctrl+B keystroke will not suspend the foreground process, but rather move one character backward in some applications. The jobs keystroke will list all processes associated with the current terminal. The bg & keystroke will cause an error because bg does not take any arguments. References: [How to Run Linux Processes in Background]

NEW QUESTION 151
A DevOps engineer is working on a local copy of a Git repository. The engineer would like to switch from the main branch to the staging branch but notices the staging branch does not exist. Which of the following Git commands should the engineer use to perform this task?

A. git branch —m staging

B. git commit —m staging
C. git status —b staging
D. git checkout —b staging

Answer: D

Explanation:
The correct answer is D. git checkout -b staging
This command will create a new branch named staging and switch to it. The git checkout command is used to switch between branches or restore files from a specific branch. The - b option is used to create a new branch if it does not exist. For example, git checkout -b staging will create and switch to the staging branch.
The other options are incorrect because:
* A. git branch -m staging
This command will rename the current branch to staging, not switch to it. The git branch command is used to list, create, or delete branches. The -m option is used to rename a branch. For example, git branch -m staging will rename the current branch to staging.
* B. git commit -m staging
This command will commit the changes in the working tree to the current branch with a message of staging, not switch to it. The git commit command is used to record changes to the repository. The -m option is used to specify a commit message. For example, git commit -m staging will commit the changes with a message of staging.
* C. git status -b staging
This command will show the status of the working tree and the current branch, not switch to it. The git status command is used to show the state of the working tree and the staged changes. The -b option is used to show the name of the current branch. However, this option does not take an argument, so specifying staging after it will cause an error. References:
? Git - git-checkout Documentation
? Git Tutorial: Create a New Branch With Git Checkout
? Git Branching - Basic Branching and Merging

NEW QUESTION 156
A DevOps engineer needs to allow incoming traffic to ports in the range of 4000 to 5000 on a Linux server. Which of the following commands will enforce this rule?

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

Answer: B

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

NEW QUESTION 160
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 165
Which of the following technologies can be used as a central repository of Linux users and groups?

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

Answer: A

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

NEW QUESTION 169
A Linux administrator is troubleshooting SSH connection issues from one of the workstations.
When users attempt to log in from the workstation to a server with the IP address 104.21.75.76, they receive the following message:

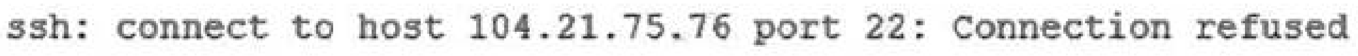

```
ssh: connect to host 104.21.75.76 port 22: Connection refused
```

The administrator reviews the information below:

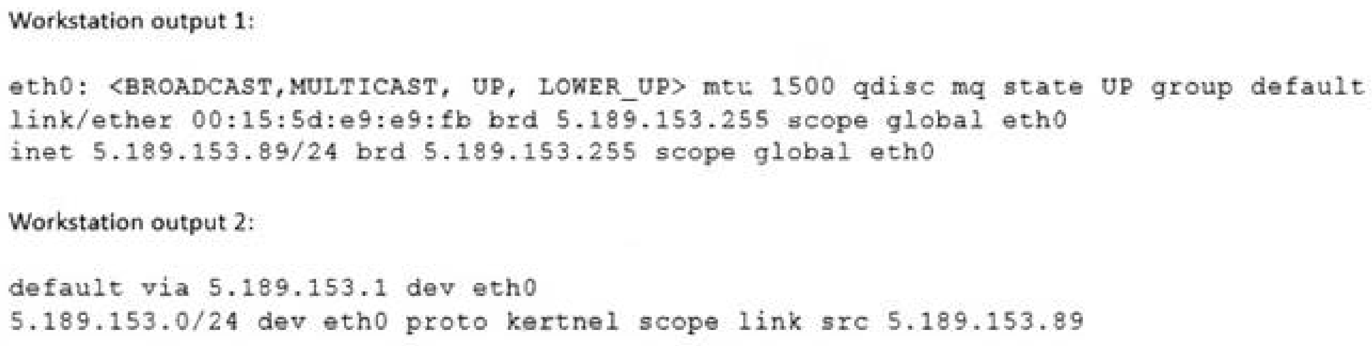

Workstation output 1:

```
eth0: <BROADCAST,MULTICAST, UP, LOWER_UP> mtu 1500 qdisc mq state UP group default
link/ether 00:15:5d:e9:e9:fb brd 5.189.153.255 scope global eth0
inet 5.189.153.89/24 brd 5.189.153.255 scope global eth0
```

Workstation output 2:

```
default via 5.189.153.1 dev eth0
5.189.153.0/24 dev eth0 proto kertnel scope link src 5.189.153.89
```

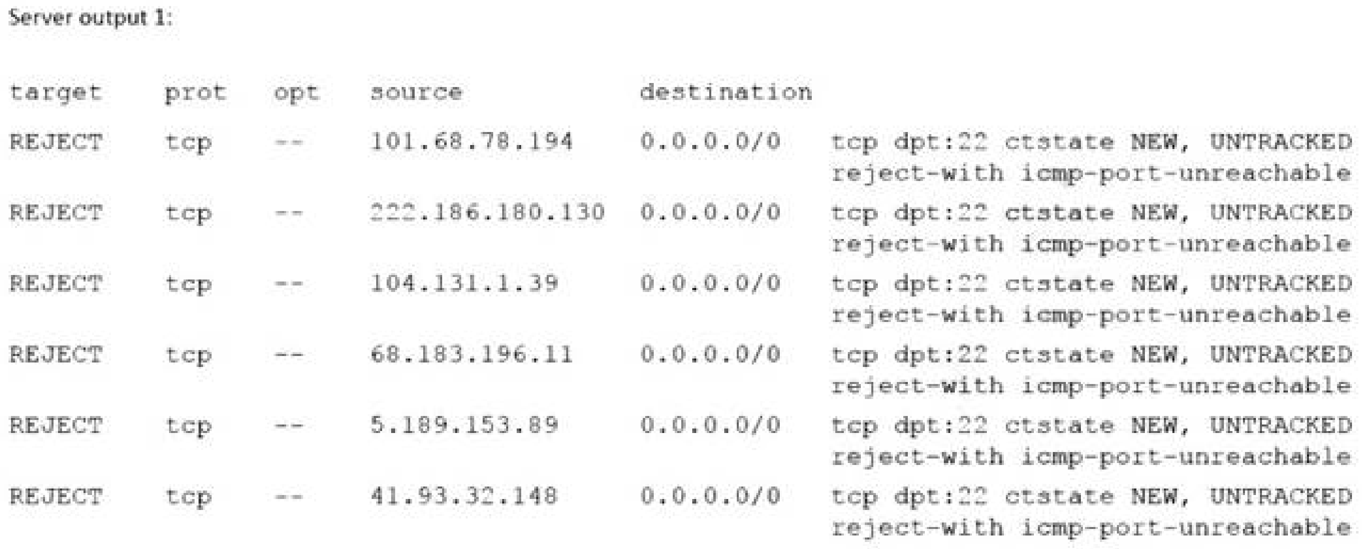

Server output 1:

```
target   prot   opt   source            destination
REJECT   tcp    --    101.68.78.194     0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                     reject-with icmp-port-unreachable
REJECT   tcp    --    222.186.180.130   0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                     reject-with icmp-port-unreachable
REJECT   tcp    --    104.131.1.39      0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                     reject-with icmp-port-unreachable
REJECT   tcp    --    68.183.196.11     0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                     reject-with icmp-port-unreachable
REJECT   tcp    --    5.189.153.89      0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                     reject-with icmp-port-unreachable
REJECT   tcp    --    41.93.32.148      0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                     reject-with icmp-port-unreachable
```

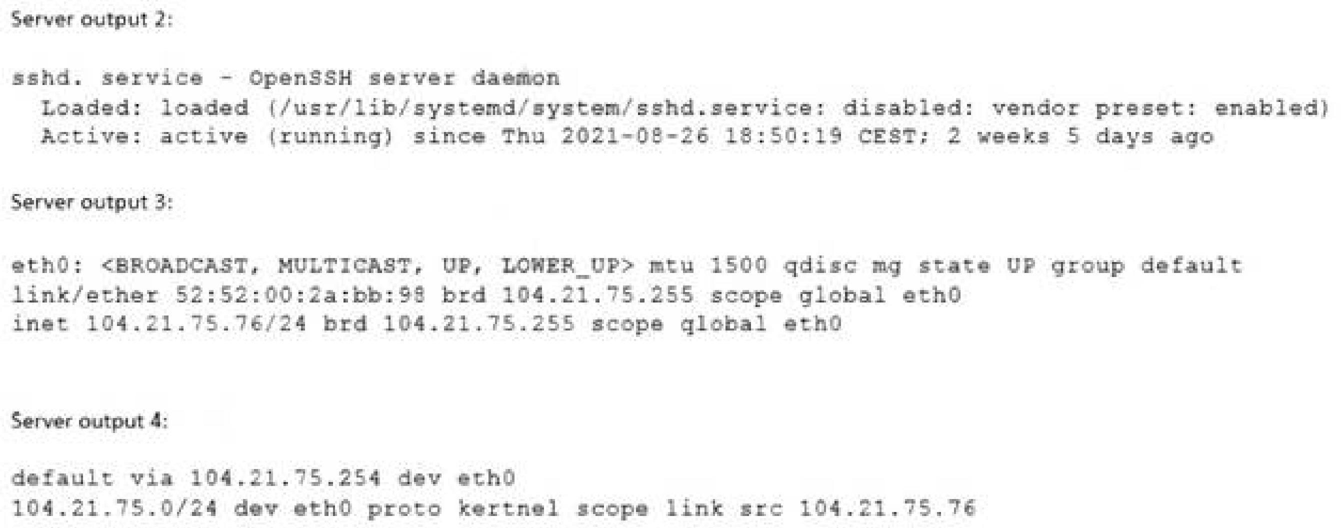

Server output 2:

```
sshd. service - OpenSSH server daemon
  Loaded: loaded (/usr/lib/systemd/system/sshd.service: disabled: vendor preset: enabled)
  Active: active (running) since Thu 2021-08-26 18:50:19 CEST; 2 weeks 5 days ago
```

Server output 3:

```
eth0: <BROADCAST, MULTICAST, UP, LOWER_UP> mtu 1500 qdisc mg state UP group default
link/ether 52:52:00:2a:bb:98 brd 104.21.75.255 scope global eth0
inet 104.21.75.76/24 brd 104.21.75.255 scope global eth0
```

Server output 4:

```
default via 104.21.75.254 dev eth0
104.21.75.0/24 dev eth0 proto kertnel scope link src 104.21.75.76
```

Which of the following is causing the connectivity issue?

A. The workstation has the wrong IP settings.
B. The sshd service is disabled.
C. The server’s firewall is preventing connections from being made.
D. The server has an incorrect default gateway configuration.

Answer: C

Explanation:
The server’s firewall is preventing connections from being made, which is causing the connectivity issue. The output of iptables -L -n shows that the firewall is blocking all incoming traffic on port 22, which is the default port for SSH. The output of ssh -v user@104.21.75.76 shows that the connection is refused by the server. To resolve the issue, the administrator needs to allow port 22 on the firewall. The other options are incorrect because they are not supported by the outputs. The workstation has the correct IP settings, as shown by the output of ip addr show. The sshd service is enabled and running, as shown by the output of

systemct1 status sshd. The server has the correct default gateway configuration, as shown by the output of ip route show. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, pages 406-407.

NEW QUESTION 173
A systems administrator needs to verify whether the built container has the app.go file in its root directory. Which of the following can the administrator use to verify the root directory has this file?

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

Answer: C

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

NEW QUESTION 175
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 180
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 182
A Linux engineer receives reports that files created within a certain group are being modified by users who are not group members. The engineer wants to reconfigure the server so that only file owners and group members can modify new files by default. Which of the following commands would accomplish this task?

A. chmod 775
B. umas
C. 002
D. chactr -Rv
E. chown -cf

Answer: B

Explanation:
The command umask 002 will accomplish the task of reconfiguring the server so that only file owners and group members can modify new files by default.
The umask command is a tool for setting the default permissions for new files and directories on Linux systems. The umask value is a four-digit octal number that represents the permissions that are subtracted from the default permissions. The default permissions for files are 666, which means read and write for owner, group, and others. The default permissions for directories are 777, which means read, write, and execute for owner, group, and others. The umask value consists

of four digits: the first digit is for special permissions, such as setuid, setgid, and sticky bit; the second digit is for the owner permissions; the third digit is for the group permissions; and the fourth digit is for the others permissions. The umask value can be calculated by subtracting the desired permissions from the default permissions. For example, if the desired permissions for files are 664, which means read and write for owner and group, and read for others, then the umask value is 002, which is 666 - 664. The command umask 002 will set the umask value to 002, which will ensure that only file owners and group members can modify new files by default. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not set the default permissions for new files (chmod 775 or chown - cf) or do not exist (chattr -Rv). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing File Permissions and Ownership, page 349.

NEW QUESTION 183

An application developer received a file with the following content:
##This is a sample Image ## FROM ubuntu:18.04
MAINTAINER demohut@gtmail.com.hac COPY . /app
RUN make /app
CMD python /app/app.py RUN apt-get update
RUN apt-get install -y nginx CMD ["echo","Image created"]
The developer must use this information to create a test bed environment and identify the image (myimage) as the first version for testing a new application before moving it to production. Which of the following commands will accomplish this task?

A. docker build -t myimage:1.0 .
B. docker build -t myimage: .
C. docker build -t myimage-1.0 .
D. docker build -i myimage:1.0 .

Answer: A

Explanation:
The docker build command is used to build an image from a Dockerfile and a context1. The Dockerfile is a text file that contains the instructions for creating the image, and the context is a set of files that can be used in the image creation process1. The file that the developer received is an example of a Dockerfile.
The -t option is used to specify a name and an optional tag for the image1. The name and tag are separated by a colon (:), and the tag is usually used to indicate the version of the image2. For example, -t myimage:1.0 means that the image will be named myimage and tagged as 1.0.
The last argument of the docker build command is the path to the context, which can be a local directory or a URL1. The dot (.) means that the current working directory is the context2. Therefore, docker build -t myimage:1.0 . means that the image will be built from the Dockerfile and the files in the current working directory, and it will be named myimage and tagged as 1.0.

NEW QUESTION 186

A Linux administrator needs to create a symlink for /usr/local/bin/app-a, which was installed in /usr/local/share/app-a. Which of the following commands should the administrator use?

A. In -s /usr/local/bin/app-a /usr/local/share/app-a
B. mv -f /usr/local/share/app-a /usr/local/bin/app-a
C. cp -f /usr/local/share/app-a /usr/local/bin/app-a
D. rsync -a /usr/local/share/app-a /usr/local/bin/app-a

Answer: A

Explanation:
To create a symlink for /usr/local/bin/app-a, which was installed in /usr/local/share/app-a, the administrator can use the command ln -s /usr/local/share/app-a /usr/local/bin/app-a (A). This will create a symbolic link named /usr/local/bin/app-a that points to the original file /usr/local/share/app-a. The other commands will not create a symlink, but either move, copy, or synchronize the file. References:
? [CompTIA Linux+ Study Guide], Chapter 3: Working with Files, Section: Creating Links
? [How to Create Symbolic Links in Linux]

NEW QUESTION 187

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:

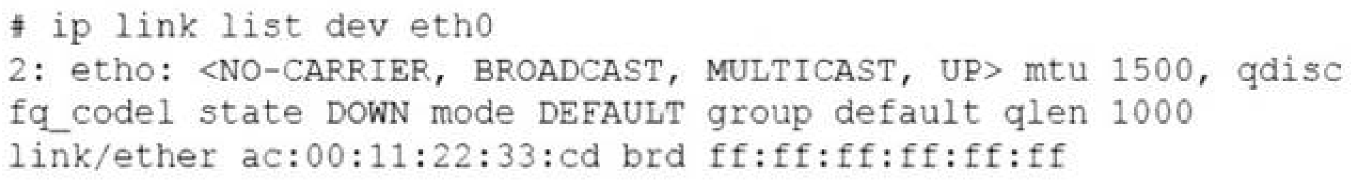

```
# ip link list dev eth0
2: etho: <NO-CARRIER, BROADCAST, MULTICAST, UP> mtu 1500, qdisc
fq_codel state DOWN mode DEFAULT group default qlen 1000
link/ether ac:00:11:22:33:cd brd ff:ff:ff:ff:ff:ff
```

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 189
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:

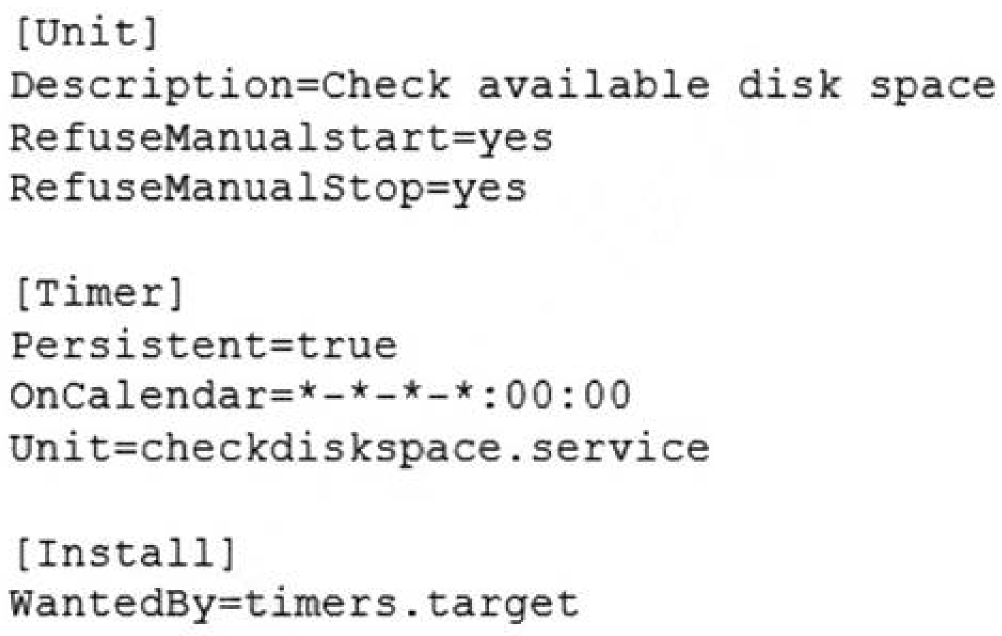

```
[Unit]
Description=Check available disk space
RefuseManualstart=yes
RefuseManualStop=yes

[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:

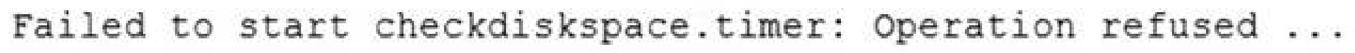

```
Failed to start checkdiskspace.timer: Operation refused ...
```

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 192
A cloud engineer wants to delete all unused networks that are not referenced by any container. Which of the following commands will achieve this goal?

A. docker network erase
B. docker network clear
C. docker network prune
D. docker network rm

Answer: C

Explanation:
The docker command is used to manage Docker containers, images, networks, volumes, and other resources on a Linux system. Docker is a platform that allows users to run applications in isolated environments called containers. Docker also provides networking features that allow users to create and manage networks for containers.
To delete all unused networks that are not referenced by any container, the cloud engineer can use the docker network prune command. This command will remove all networks that have no containers connected to them. The statement C is correct.
The statements A, B, and D are incorrect because they do not delete all unused networks.
The docker network erase and docker network clear commands do not exist. The docker network rm command deletes a specific network by name or ID, but not all unused networks. References: [How to Manage Docker Networks]

NEW QUESTION 194
Joe, a user, is unable to log in to the Linux system. Given the following output:

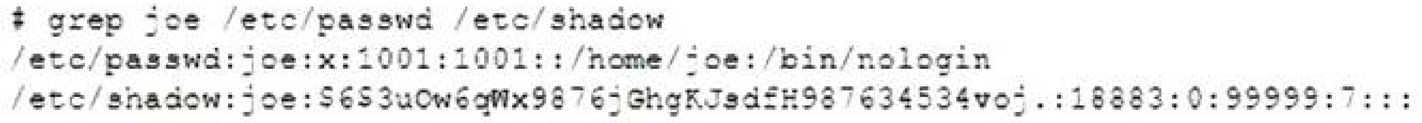

```
# grep joe /etc/passwd /etc/shadow
/etc/passwd:joe:x:1001:1001::/home/joe:/bin/nologin
/etc/shadow:joe:$6$3uOw6qWx9876jGhgKJadfH987634534voj.:18883:0:99999:7:::
```

Which of the following commands would resolve the issue?

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

Answer: B

Explanation:
The command pam_tally2 -u joe -r will resolve the issue of Joe being unable to log in to the Linux system. The pam_tally2 command is a tool for managing the login counter for the PAM (Pluggable Authentication Modules) system. PAM is a framework for managing authentication and authorization on Linux systems. PAM allows the administrator to define the rules and policies for accessing various system resources and services, such as login, sudo, ssh, or cron. PAM also supports different types of authentication methods, such as passwords, tokens, biometrics, or smart cards. PAM can be used to implement login restrictions, such as limiting the number of failed login attempts, locking the account after a certain number of failures, or enforcing a minimum or maximum time between login attempts. The pam_tally2 command can display, reset, or unlock the login counter for the users or hosts. The -u joe option specifies the user name that the command should apply to. The -r option resets the login counter for the user. The command pam_tally2 -u joe - r will reset the login counter for Joe, which will unlock his account and allow him to log in to the Linux system. This will resolve the issue of Joe being unable to log in to the Linux system. This is the correct command to use to resolve the issue. The other options are incorrect because they either do not unlock the account (usermod -s /bin/bash joe or passwd -u joe) or do not affect the login counter (chage -E 90
joe). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Implementing Basic Security, page 517.

NEW QUESTION 198
Which of the following commands is used to configure the default permissions for new files?

A. setenforce
B. sudo
C. umask
D. chmod

Answer: C

Explanation:
The command that is used to configure the default permissions for new files is umask. The umask command is a tool for setting the default permissions for new files and directories on Linux systems. The umask value is a four-digit octal number that represents the permissions that are subtracted from the default permissions. The default permissions for files are 666, which means read and write for owner, group, and others.
The default permissions for directories are 777, which means read, write, and execute for owner, group, and others. The umask value consists of four digits: the first digit is for special permissions, such as setuid, setgid, and sticky bit; the second digit is for the owner permissions; the third digit is for the group permissions; and the fourth digit is for the others permissions. The umask value can be calculated by subtracting the desired permissions from the default permissions. For example, if the desired permissions for files are 664, which means read and write for owner and group, and read for others, then the umask
value is 002, which is 666 - 664. The command umask 002 will set the umask value to 002, which will ensure that only file owners and group members can modify new files by default. The command that is used to configure the default permissions for new files is umask. This is the correct answer to the question. The other options are incorrect because they either do not set the default permissions for new files (setenforce, sudo, or chmod) or do not exist (kill -HUP or kill -TERM).
References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing File Permissions and Ownership, page 349.

NEW QUESTION 200
As part of the requirements for installing a new application, the swappiness parameter needs to be changed to O. This change needs to persist across re-boots and be applied immediately. A Linux systems administrator is performing this change. Which of the following steps should the administrator complete to accomplish this task?

A. echo "v
B. swappiness—()" >> /etc/sysctl . conf && sysctl —p
C. echo "vr
D. >> / proc/meminfo && sysctl —a
E. sysctl —v >> / proc/meminfo & & echo "v
F. swapiness=0"
G. sysctl —h "v
H. swapiness—O" && echo / etc/vmswapiness

Answer: A

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

NEW QUESTION 201
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 205

Ann, a security administrator, is performing home directory audits on a Linux server. Ann issues the su Joe command and then issues the Is command. The output displays files that reside in Ann's home directory instead of Joe's. Which of the following represents the command Ann should have issued in order to list Joe's files?

A. su - Joe
B. sudo Joe
C. visudo Joe
D. pkexec joe

Answer: A

Explanation:

The su command is used to switch to another user account on Linux systems. The - option makes the shell a login shell, which means that it will read the profile and environment variables of the target user. Without this option, the shell will retain the environment variables of the original user. This can cause confusion when issuing commands that depend on these variables, such as ls, which uses the $HOME variable to determine the home directory. Therefore, Ann should have issued su - Joe to list Joe’s files instead of her own. References: [How to Use su Command in Linux with Examples]

NEW QUESTION 210

A systems administrator detected corruption in the /data filesystem. Given the following output:

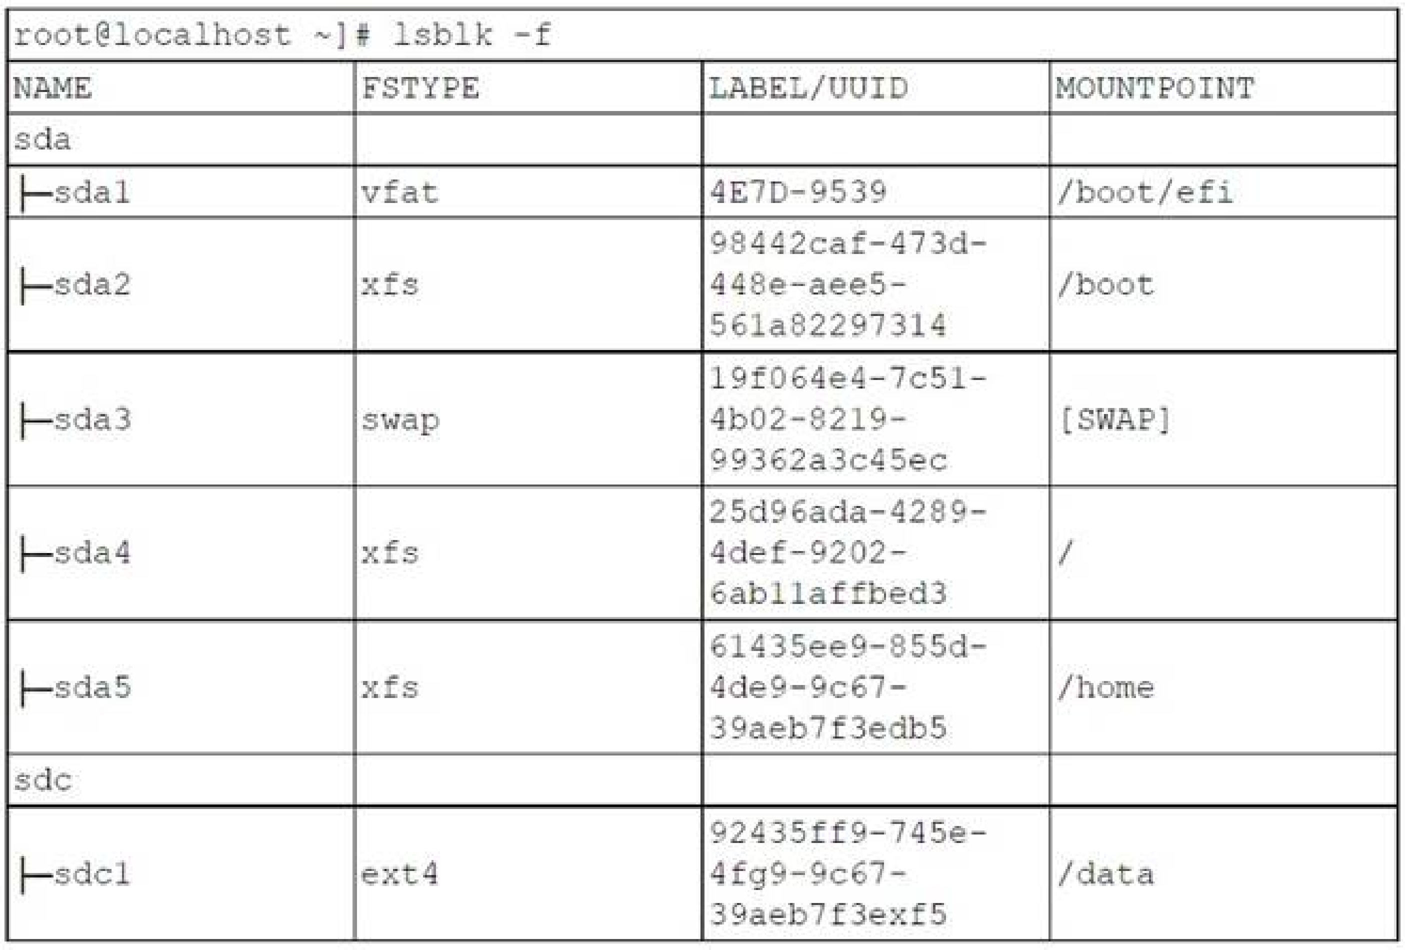

```
root@localhost ~]# lsblk -f
```

NAME	FSTYPE	LABEL/UUID	MOUNTPOINT
sda			
├sda1	vfat	4E7D-9539	/boot/efi
├sda2	xfs	98442caf-473d-448e-aee5-561a82297314	/boot
├sda3	swap	19f064e4-7c51-4b02-8219-99362a3c45ec	[SWAP]
├sda4	xfs	25d96ada-4289-4def-9202-6ab11affbed3	/
├sda5	xfs	61435ee9-855d-4de9-9c67-39aeb7f3edb5	/home
sdc			
├sdc1	ext4	92435ff9-745e-4fg9-9c67-39aeb7f3exf5	/data

Which of the following commands can the administrator use to best address this issue?

A. umount /data mkfs . xfs /dev/sclcl mount /data
B. umount /data xfs repair /dev/ sdcl mount /data
C. umount /data fsck /dev/ sdcl mount / data
D. umount /data pvs /dev/sdcl mount /data

Answer: B

Explanation:

The xfs repair command is used to check and repair an XFS filesystem, which is the type of filesystem used for the /data partition, as shown in the output. The administrator needs to unmount the /data partition before running the xfs repair command on it, and then mount it back after the repair is done. For example: umount /data; xfs_repair /dev/sdcl; mount /data. The mkfs.xfs command is used to create a new XFS filesystem, which would erase all the data on the partition. The fsck command is used to check and repair other types of filesystems, such as ext4, but not XFS. The pvs command is used to display information about physical volumes in a logical volume manager (LVM) setup, which is not relevant for this issue.

NEW QUESTION 213

SIMULATION

Junior system administrator had trouble installing and running an Apache web server on a Linux server. You have been tasked with installing the Apache web server on the Linux server and resolving the issue that prevented the junior administrator from running Apache.

INSTRUCTIONS

Install Apache and start the service. Verify that the Apache service is running with the defaults.

Typing “help” in the terminal will show a list of relevant event commands.

If at any time you would like to bring back the initial state of the simulation, please click the Reset All button.

CentOS Command Prompt

```
[root@centos7]#
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
yum install httpd
systemct1 --now enable httpd systemct1 status httpd netstat -tunlp | grep 80
pkill <processname> systemct1 restart httpd systemct1 status httpd

NEW QUESTION 214
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 219
A systems administrator is notified that the mysqld process stopped unexpectedly. The systems administrator issues the following command: sudo grep –i -r ‘out of memory’ /var/log
The output of the command shows the following:
kernel: Out of memory: Kill process 9112 (mysqld) score 511 or sacrifice child.
Which of the following commands should the systems administrator execute NEXT to troubleshoot this issue? (Select two).

A. free -h
B. nc -v 127.0.0.1 3306
C. renice -15 $(pidof mysql)
D. lsblk
E. killall -15
F. vmstat -a 1 4

Answer: AF

Explanation:
The free -h command can be used to check the amount of free and used memory in the system in a human-readable format. This can help to troubleshoot the issue of mysqld being killed due to out of memory. The vmstat -a 1 4 command can be used to monitor the system’s virtual memory statistics, such as swap usage, paging activity, and memory faults, every one second for four times. This can help to identify any memory pressure or performance issues that may cause out of memory errors. The nc -v 127.0.0.1 3306 command would attempt to connect to the MySQL server on port 3306 and display any diagnostic messages, but

this would not help to troubleshoot the memory issue. The renice -15 $(pidof mysql) command would change the priority of the mysql process to -15, but this would not prevent it from being killed due to out of memory. The lsblk command would display information about block devices, not memory usage. The killall -15 command would send a SIGTERM signal to all processes with a matching name, but this would not help to troubleshoot the memory issue. References: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 15: Managing Memory and Process Execution, pages 468-469.

NEW QUESTION 224

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 225

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 228

......