Exam Questions XK0-005

CompTIA Linux+ Certification Exam

https://www.2passeasy.com/dumps/XK0-005/

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

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

Answer: D

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

NEW QUESTION 2
The administrator comptia is not able to perform privileged functions on a newly deployed system. Given the following command outputs:

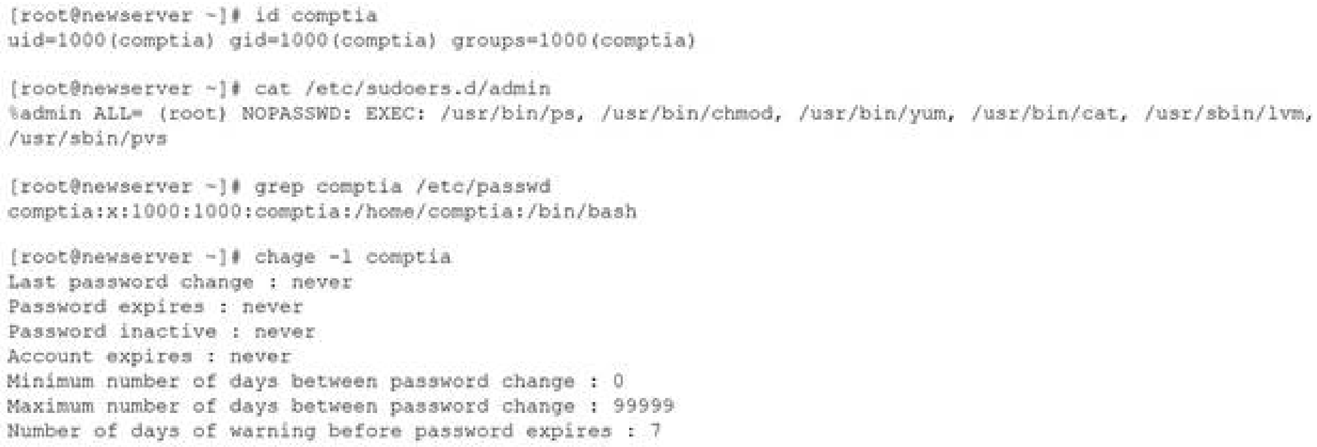

```
[root@newserver ~]# id comptia
uid=1000(comptia) gid=1000(comptia) groups=1000(comptia)

[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 3
A Linux administrator is tasked with creating resources using containerization. When deciding how to create this type of deployment, the administrator identifies some key features, including portability, high availability, and scalability in production. Which of the following should the Linux administrator choose for the new design?

A. Docker
B. On-premises systems
C. Cloud-based systems
D. Kubernetes

Answer: D

Explanation:
The Linux administrator should choose Kubernetes for the new design that requires portability, high availability, and scalability in production using containerization. Kubernetes is an open-source platform that automates the deployment, scaling, and management of containerized applications across clusters of nodes. Kubernetes provides features such as service discovery, load balancing, storage orchestration, self-healing, secret and configuration management, and batch execution. Kubernetes also supports multiple container runtimes, such as Docker, containerd, and CRI-O, making it portable across different platforms and clouds. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks; What is Kubernetes? | Kubernetes

NEW QUESTION 4

A Linux administrator has physically added a new RAID adapter to a system. Which of the following commands should the Linux administrator run to confirm that the device has been recognized? (Select TWO).

A. rmmod
B. Is -11 /etc
C. Ishw —class disk
D. pvdisplay
E. rmdir /dev
F. dmesg

Answer: CF

Explanation:
The following commands can help you confirm that the new RAID adapter has been recognized by the Linux system:
? dmesg: This command displays the kernel messages, which can show the information about the newly detected hardware device. You can use dmesg | grep -i raid to filter the output for RAID-related messages.
? lshw -class disk: This command lists the disk devices on the system, including the RAID controller and its model name. You can use lshw -class disk | grep -i raid to filter the output for RAID-related information1.
The other commands are not relevant for this purpose. For example:
? rmmod: This command removes a module from the Linux kernel, which is not useful for detecting a new device.
? ls -l /etc: This command lists the files and directories in the /etc directory, which is not related to hardware devices.
? pvdisplay: This command displays the attributes of physical volumes, which are part of the logical volume management (LVM) system, not the RAID system.
? rmdir /dev: This command removes an empty directory, which is not helpful for detecting a new device. Moreover, /dev is a special directory that contains device files, and should not be removed.

NEW QUESTION 5
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 6
The application team has reported latency issues that are causing the application to crash on the Linux server. The Linux administrator starts troubleshooting and receives the following output:

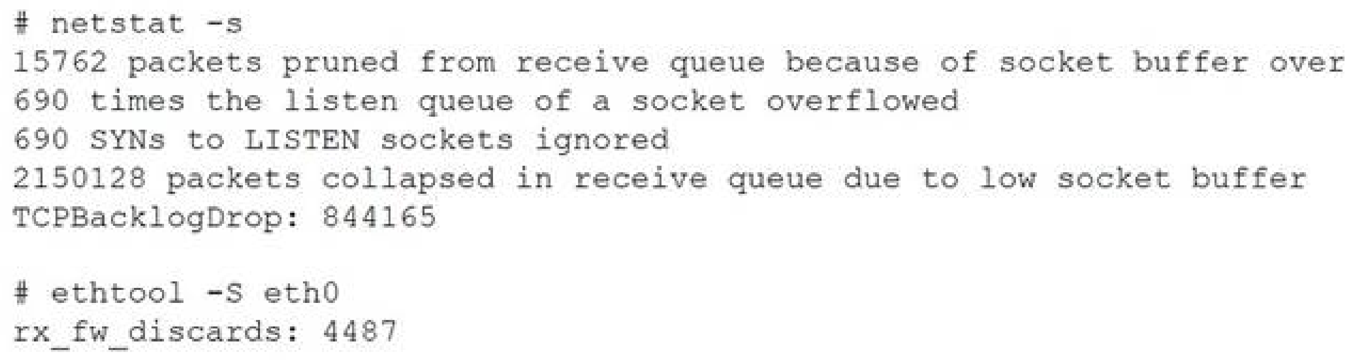

```
# netstat -s
15762 packets pruned from receive queue because of socket buffer over
690 times the listen queue of a socket overflowed
690 SYNs to LISTEN sockets ignored
2150128 packets collapsed in receive queue due to low socket buffer
TCPBacklogDrop: 844165

# ethtool -S eth0
rx_fw_discards: 4487
```

Which of the following commands will improve the latency issue?

A. # echo 'net.core.net_backlog = 5000000' >> /etc/sysctl.conf# sysctl -p# systemctl daemon-reload
B. # ifdown eth0# ip link set dev eth0 mtu 800# ifup eth0
C. # systemctl stop network# ethtool -g eth0 512# systemctl start network
D. # echo 'net.core.rmem max = 12500000' >> /etc/sysctl.conf# echo 'net.core.wmem_max = 12500000' >> /etc/sysctl.conf# sysctl -p

Answer: D

Explanation:
The best command to use to improve the latency issue is D. # echo ‘net.core.rmem max = 12500000’ >> /etc/sysctl.conf # echo ‘net.core.wmem_max = 12500000’ >> /etc/sysctl.conf # sysctl -p. This command will increase the size of the receive and send buffers for the network interface, which can improve the network performance and reduce packet loss. The sysctl command will apply the changes to the kernel parameters without rebooting the system.
The other commands are either incorrect or not suitable for this task. For example:
? A. # echo ‘net.core.net_backlog = 5000000’ >> /etc/sysctl.conf # sysctl -p # systemctl daemon-reload will try to increase the backlog queue for incoming connections, but this is not relevant for the latency issue. The systemctl daemon- reload command is also unnecessary, as it only reloads the systemd configuration files, not the kernel parameters.
? B. # ifdown eth0 # ip link set dev eth0 mtu 800 # ifup eth0 will try to change the maximum transmission unit (MTU) of the network interface to 800 bytes, but this is too low and may cause fragmentation and performance degradation. The default MTU for Ethernet is 1500 bytes, and it should not be changed unless there is a

specific reason.
? C. # systemctl stop network # ethtool -g eth0 512 # systemctl start network will try to change the ring buffer size of the network interface to 512, but this is too small and may cause packet drops and latency spikes. The default ring buffer size for Ethernet is usually 4096 or higher, and it should be increased if there is a high network traffic.

NEW QUESTION 7
A Linux administrator is troubleshooting a memory-related issue. Based on the output of the commands:

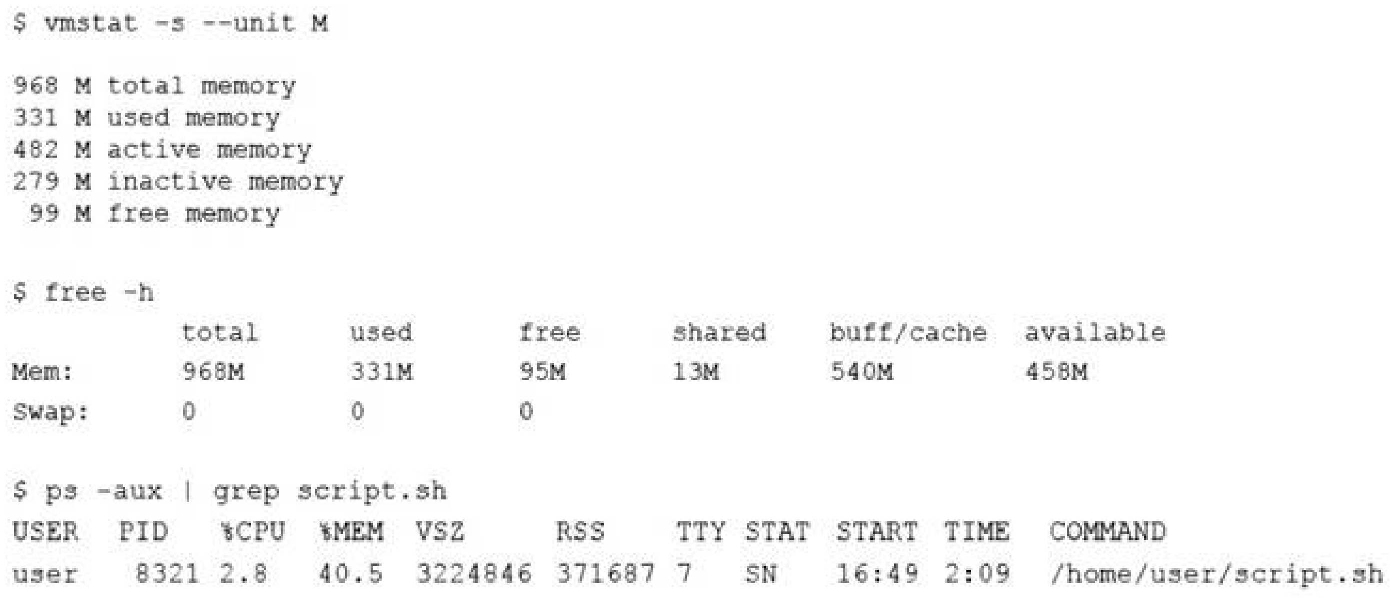

```
$ vmstat -s --unit M

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 8
In order to copy data from another VLAN, a systems administrator wants to temporarily assign IP address 10.0.6 5/24 to the newly added network interface enp1s0f1. Which of the following commands should the administrator run to achieve the goal?

A. ip addr add 10.0.6.5/24 dev enpls0f1
B. echo "IPV4_ADDRESS=10.0.6.5/24" > /etc/sysconfig/network-scripts/ifcfg-enplsOfl
C. ifconfig 10.0.6.5/24 enpsIs0f1
D. nmcli conn add lpv4.address-10.0.6.5/24 ifname enpls0f1

Answer: A

Explanation:
The command ip addr add 10.0.6.5/24 dev enp1s0f1 will achieve the goal of temporarily assigning IP address 10.0.6.5/24 to the newly added network interface enp1s0f1. The ip command is a tool for managing network interfaces and routing on Linux systems. The addr option specifies the address manipulation mode. The add option adds a new address to an interface. The 10.0.6.5/24 is the IP address and the subnet mask in CIDR notation. The dev option specifies the device name. The enp1s0f1 is the name of the network interface. The command ip addr add 10.0.6.5/24 dev enp1s0f1 will add the IP address 10.0.6.5/24 to the network interface enp1s0f1, which will allow the administrator to copy data from another VLAN. This is the correct command to use to achieve the goal. The other options are incorrect because they either do not add a new address to an interface (echo "IPV4_ADDRESS=10.0.6.5/24" > /etc/sysconfig/network-scripts/ifcfg- enp1s0f1 or ifconfig 10.0.6.5/24 enp1s0f1) or do not use the correct syntax for the command (nmcli conn add ipv4.address-10.0.6.5/24 ifname enp1s0f1 instead of nmcli conn add type ethernet ipv4.address 10.0.6.5/24 ifname enp1s0f1). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 385.

NEW QUESTION 9
Some servers in an organization have been compromised. Users are unable to access to the organization’s web page and other services. While reviewing the system log, a systems administrator notices messages from the kernel regarding firewall rules:

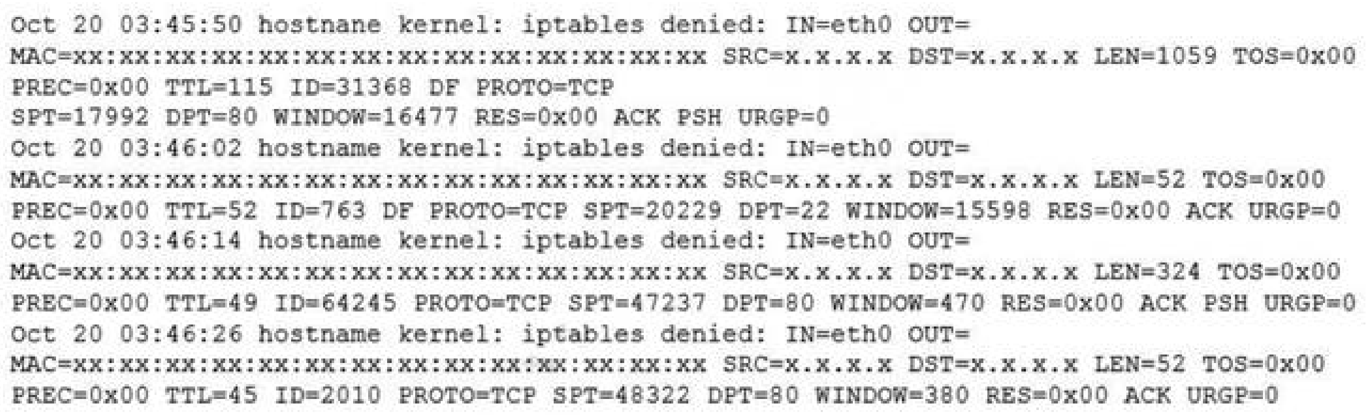

```
Oct 20 03:45:50 hostnane kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx SRC=x.x.x.x DST=x.x.x.x LEN=1059 TOS=0x00
PREC=0x00 TTL=115 ID=31368 DF PROTO=TCP
SPT=17992 DPT=80 WINDOW=16477 RES=0x00 ACK PSH URGP=0
Oct 20 03:46:02 hostname kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx SRC=x.x.x.x DST=x.x.x.x LEN=52 TOS=0x00
PREC=0x00 TTL=52 ID=763 DF PROTO=TCP SPT=20229 DPT=22 WINDOW=15598 RES=0x00 ACK URGP=0
Oct 20 03:46:14 hostname kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx SRC=x.x.x.x DST=x.x.x.x LEN=324 TOS=0x00
PREC=0x00 TTL=49 ID=64245 PROTO=TCP SPT=47237 DPT=80 WINDOW=470 RES=0x00 ACK PSH URGP=0
Oct 20 03:46:26 hostname kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx SRC=x.x.x.x DST=x.x.x.x LEN=52 TOS=0x00
PREC=0x00 TTL=45 ID=2010 PROTO=TCP SPT=48322 DPT=80 WINDOW=380 RES=0x00 ACK URGP=0
```

Which of the following commands will remediate and help resolve the issue?

A.

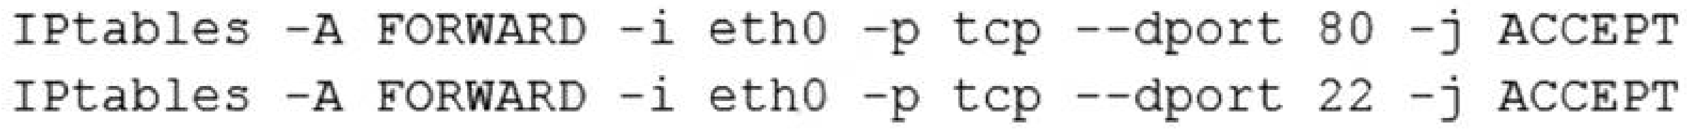

```
IPtables -A FORWARD -i eth0 -p tcp --dport 80 -j ACCEPT
IPtables -A FORWARD -i eth0 -p tcp --dport 22 -j ACCEPT
```

B.

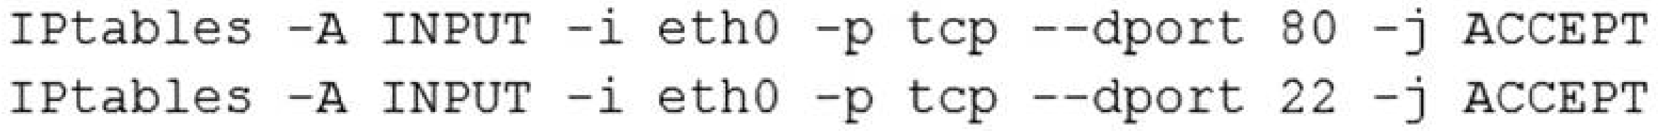

```
IPtables -A INPUT -i eth0 -p tcp --dport 80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --dport 22 -j ACCEPT
```

C.

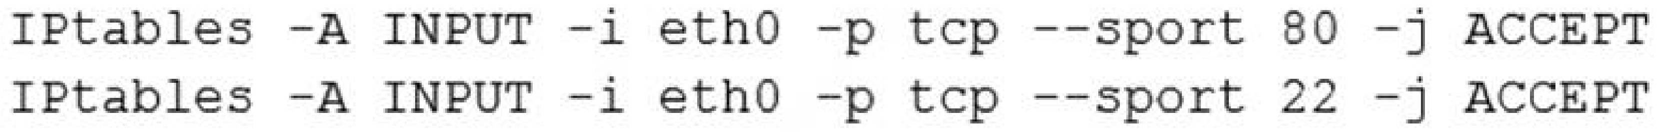

```
IPtables -A INPUT -i eth0 -p tcp --sport 80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --sport 22 -j ACCEPT
```

D.

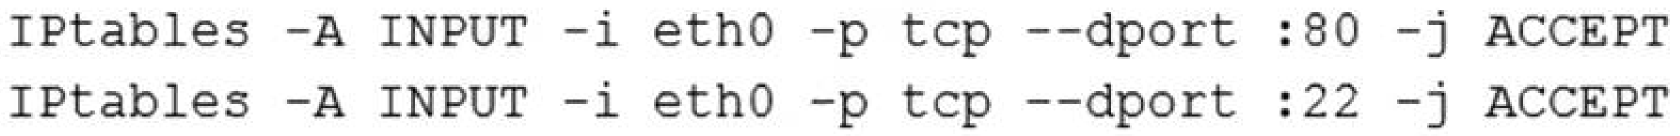

```
IPtables -A INPUT -i eth0 -p tcp --dport :80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --dport :22 -j ACCEPT
```

Answer: A

Explanation:
The command iptables -F will remediate and help resolve the issue. The issue is caused by the firewall rules that block the access to the organization’s web page and other services. The output of dmesg | grep firewall shows that the kernel has dropped packets from the source IP address 192.168.1.100 to the destination port 80, which is the default port for HTTP. The command iptables -F will flush all the firewall rules and allow the traffic to pass through. This command will resolve the issue and restore the access to the web page and other services. The other options are incorrect because they either do not affect the firewall rules (ip route flush or ip addr flush) or do not exist (iptables - R). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 543.

NEW QUESTION 10
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 10
User1 is a member of the accounting group. Members of this group need to be able to execute but not make changes to a script maintained by User2. The script should not be accessible to other users or groups. Which of the following will give proper access to the script?

A. chown user2:accounting script.sh chmod 750 script.sh
B. chown user1:accounting script.shchmod 777 script.sh
C. chown accounting:user1 script.sh chmod 057 script.sh

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

Answer: A

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

NEW QUESTION 15
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 20
An administrator runs ping comptia.org. The result of the command is:
ping: comptia.org: Name or service not known
Which of the following files should the administrator verify?

A. /etc/ethers
B. /etc/services
C. /etc/resolv.conf
D. /etc/sysctl.conf

Answer: C

Explanation:
The best file to verify when the ping command returns the error “Name or service not known” is C. /etc/resolv.conf. This file contains the configuration for the DNS resolver, which is responsible for translating domain names into IP addresses. If this file is missing, corrupted, or has incorrect entries, the ping command will not be able to resolve the domain name and will fail with the error. To fix this issue, the administrator should check that the file exists, has proper permissions, and has valid nameserver entries. For example, a typical /etc/resolv.conf file may look like this:
nameserver 8.8.8.8 nameserver 8.8.4.4
These are the IP addresses of Google’s public DNS servers, which can be used as a fallback option if the default DNS servers are not working.

NEW QUESTION 24
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 28
A Linux administrator wants to prevent the httpd web service from being started both manually and automatically on a server. Which of the following should the

administrator use to accomplish this task?

A. systemctl mask httpd
B. systemctl disable httpd
C. systemctl stop httpd
D. systemctl reload httpd

Answer: A

Explanation:
The best command to use to prevent the httpd web service from being started both manually and automatically on a server is A. systemctl mask httpd. This command will create a symbolic link from the httpd service unit file to /dev/null, which will make the service impossible to start or enable. This is different from systemctl disable httpd, which will only prevent the service from starting automatically on boot, but not manually. The other commands are either not relevant or not sufficient for this task. For example:
? C. systemctl stop httpd will only stop the service if it is currently running, but it will not prevent it from being started again.
? D. systemctl reload httpd will only reload the configuration files of the service, but it
will not stop or disable it.

NEW QUESTION 33
A Linux systems administrator needs to copy files and directories from Server A to Server

A. Which of the following commands can be used for this purpose? (Select TWO)
B. rsyslog
C. cp
D. rsync
E. reposync
F. scp
G. ssh

Answer: CE

Explanation:
The rsync and scp commands can be used to copy files and directories from Server A to Server B. Both commands can use SSH as a secure protocol to transfer data over the network. The rsync command can synchronize files and directories between two locations, using various options to control the copying behavior. The scp command can copy files and directories between two hosts, using similar syntax as cp. The rsyslog command is used to manage system logging, not file copying. The cp command is used to copy files and directories within a single host, not between two hosts. The reposync command is used to synchronize a remote yum repository to a local directory, not copy files and directories between two hosts. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Networking Fundamentals, pages 440-441.

NEW QUESTION 37
Which of the following directories is the mount point in a UEFI system?

A. /sys/efi
B. /boot/efi
C. /efi
D. /etc/efi

Answer: B

Explanation:
The /boot/efi directory is the mount point in a UEFI system. This directory contains the EFI System Partition (ESP), which stores boot loaders and other files required by UEFI firmware. The /sys/efi directory does not exist by default in Linux systems. The /efi directory does not exist by default in Linux systems. The /etc/efi directory does not exist by default in Linux systems. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing the Linux Boot Process, page 398.

NEW QUESTION 38
A Linux administrator is installing a web server and needs to check whether web traffic has already been allowed through the firewall. Which of the following commands should the administrator use to accomplish this task?

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

Answer: C

Explanation:
The command firewall-cmd --query-service http will accomplish the task of checking whether web traffic has already been allowed through the firewall. The firewall- cmd command is a tool for managing firewalld, which is a firewall service that provides dynamic and persistent network security on Linux systems. The firewalld uses zones and services to define the rules and policies for the network traffic. The zones are logical groups of network interfaces and sources that have the same level of trust and security. The services are predefined sets of ports and protocols that are associated with certain applications or functions. The --query-service http option queries whether a service is enabled in a zone. The http is the name of the service that the command should check.
The http service represents the web traffic that uses the port 80 and the TCP protocol. The command firewall-cmd --query-service http will check whether the http service is enabled in the default zone, which is usually the public zone. The command will return yes if the web traffic has already been allowed through the firewall, or no if the web traffic has not been allowed through the firewall. This is the correct command to use to accomplish the task.
The other options are incorrect because they either do not exist (firewalld query-service- http or firewalld --check-service http) or do not query the service (firewall-cmd --check-
service http instead of firewall-cmd --query-service http). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 392.

NEW QUESTION 39
A systems administrator wants to be sure the sudo rules just added to /etc/sudoers are valid. Which of the following commands can be used for this task?

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

Answer: A

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

NEW QUESTION 40
A Linux systems administrator is troubleshooting an I/O latency on a single CPU server. The administrator runs a top command and receives the following output:
%Cpu(s): 0.2 us, 33.1 sy, 0.0 ni, 0.0 id, 52.4 wa, 0.0 hi, 0.2 si, 0.0 st
Which of the following is correct based on the output received from the exe-cuted command?

A. The server's CPU is taking too long to process users' requests.
B. The server's CPU shows a high idle-time value.
C. The server's CPU is spending too much time waiting for data inputs.
D. The server's CPU value for the time spent on system processes is low.

Answer: C

Explanation:
The server's CPU is spending too much time waiting for data inputs. This can be inferred from the output of the top command, which shows the percentage of CPU time spent in different states. The wa state stands for wait, and it indicates that the CPU is idle while waiting for an I/O operation to complete. In this case, the wa state is 52.4%, which means that more than half of the CPU time is wasted on waiting for data inputs. This can cause a high I/O latency and affect the performance of the server.
The other options are not correct based on the output received from the executed command. The server's CPU is not taking too long to process users' requests, because the us state, which stands for user, is only 0.2%, which means that the CPU is barely used by user processes. The server's CPU does not show a high idle-time value, because the id state, which stands for idle, is 0.0%, which means that the CPU is not idle at all. The server's CPU value for the time spent on system processes is not low, because the sy state, which stands for system, is 33.1%, which means that the CPU is heavily used by system processes.
References: How to Use the Linux top Command (and Understand Its Output); [Understanding Linux CPU Load - when should you be worried?]

NEW QUESTION 41
The development team wants to prevent a file from being modified by all users in a Linux system, including the root account. Which of the following commands can be used to accomplish this objective?

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

Answer: C

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

NEW QUESTION 43
A User on a Linux workstation needs to remotely start an application on a Linux server and then forward the graphical display of that application back to the Linux workstation. Which of the following would enable the user to perform this action?

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

Answer: A

Explanation:
The ssh -X option enables X11 forwarding, which allows the user to run graphical applications on the remote server and display them on the local workstation. The user needs to specify the username, the server address, and the application name after the ssh -X command. The remote server also needs to have X11Forwarding enabled and xauth installed for this to work. References:
? The web search result 8 explains how to run a GUI application through SSH by configuring both the SSH client and server.
? The web search result 6 provides a detailed answer on how to forward X over SSH to run graphics applications remotely, with examples and troubleshooting tips.

? The CompTIA Linux+ Certification Exam Objectives mention that the candidate should be able to “use SSH for remote access and management” as part of the System Operation and Maintenance domain1.

NEW QUESTION 48
An engineer needs to insert a character at the end of the current line in the vi text editor. Which of the following will allow the engineer to complete this task?

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

Answer: D

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

NEW QUESTION 50
A junior administrator is trying to set up a passwordless SSH connection to one of the servers. The administrator follows the instructions and puts the key in the authorized_key file at the server, but the administrator is still asked to provide a password during the connection.
Given the following output:

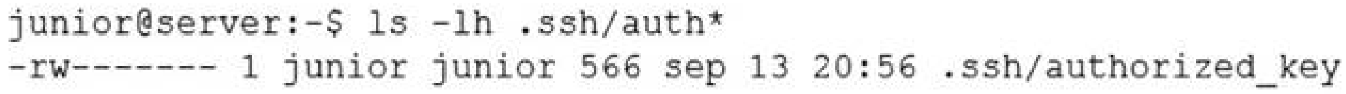

```
junior@server:-$ ls -lh .ssh/auth*
-rw------- 1 junior junior 566 sep 13 20:56 .ssh/authorized_key
```

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

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

Answer: B

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

NEW QUESTION 54
A Linux system is getting an error indicating the root filesystem is full. Which of the following commands should be used by the systems administrator to resolve this issue? (Choose three.)

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

Answer: ACE

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

rootlv print). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 318-319, 331-332.

NEW QUESTION 58
A Linux system fails to start and delivers the following error message:

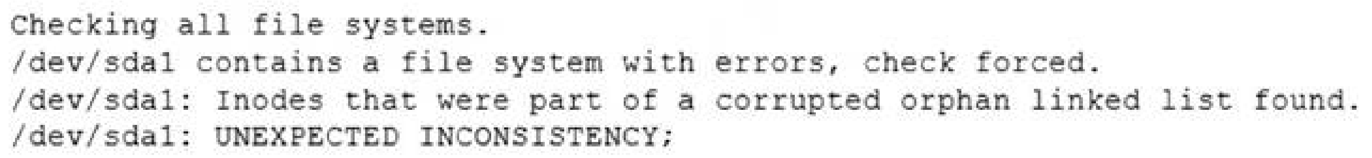

```
Checking all file systems.
/dev/sda1 contains a file system with errors, check forced.
/dev/sda1: Inodes that were part of a corrupted orphan linked list found.
/dev/sda1: UNEXPECTED INCONSISTENCY;
```

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

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

Answer: A

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

NEW QUESTION 61
A Linux administrator is trying to remove the ACL from the file /home/user/data. txt but receives the following error message:

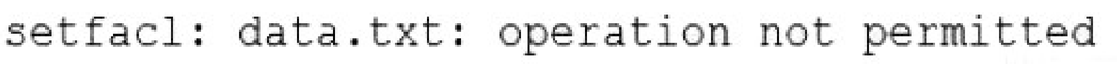

```
setfacl: data.txt: operation not permitted
```

Given the following analysis:

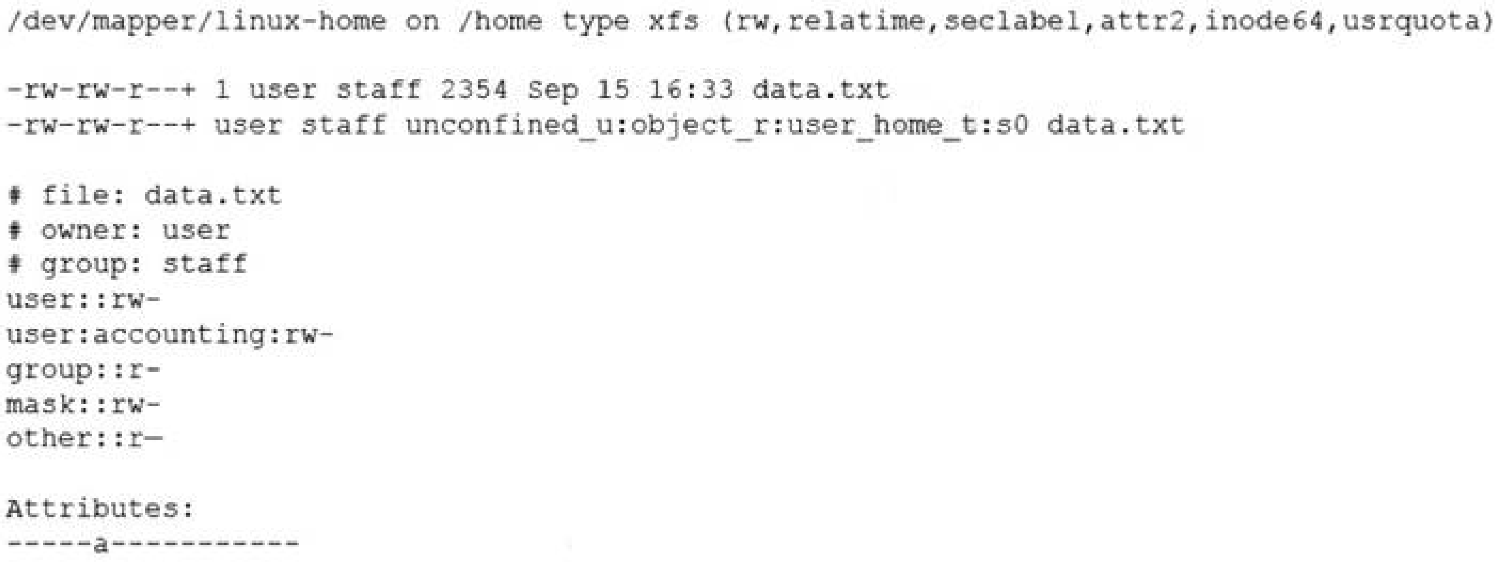

```
/dev/mapper/linux-home on /home type xfs (rw,relatime,seclabel,attr2,inode64,usrquota)

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

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

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

Which of the following is causing the error message?

A. The administrator is not using a highly privileged account.
B. The filesystem is mounted with the wrong options.
C. SELinux file context is denying the ACL changes.
D. File attributes are preventing file modification.

Answer: D

Explanation:
File attributes are preventing file modification, which is causing the error message. The output of lsattr /home/user/data.txt shows that the file has the immutable attribute (i) set, which means that the file cannot be changed, deleted, or renamed. The command setfacl -b /home/user/data.txt tries to remove the ACL from the file, but fails because of the immutable attribute. The administrator needs to remove the immutable attribute first by using the command chattr -i /home/user/data.txt and then try to remove the ACL again. The other options are incorrect because they are not supported by the outputs. The administrator is using a highly privileged account, as shown by the # prompt. The filesystem is mounted with the correct options, as shown by the output of mount | grep /home. SELinux file context is not denying the ACL changes, as shown by the output of ls - Z /home/user/data.txt. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, pages 357-358.

NEW QUESTION 66
A Linux administrator wants to set the SUID of a file named dev_team.text with 744 access rights. Which of the following commands will achieve this goal?

A. chmod 4744 dev_team.txt
B. chmod 744 --setuid dev_team.txt
C. chmod -c 744 dev_team.txt
D. chmod -v 4744 --suid dev_team.txt

Answer: A

Explanation:
The command that will set the SUID of a file named dev_team.txt with 744 access rights is chmod 4744 dev_team.txt. This command will use the chmod utility to

change the file mode bits of dev_team.txt. The first digit (4) represents the SUID bit, which means that when someone executes dev_team.txt, it will run with the permissions of the file owner. The next three digits (744) represent the read, write, and execute permissions for the owner (7), group (4), and others (4). This means that the owner can read, write, and execute dev_team.txt, while the group and others can only read it.
The other options are not correct commands for setting the SUID of a file with 744 access rights. The chmod 744 --setuid dev_team.txt command is invalid because there is no -- setuid option in chmod. The chmod -c 744 dev_team.txt command will change the file mode bits to 744, but it will not set the SUID bit. The -c option only means that chmod will report when a change is made. The chmod -v 4744 --suid dev_team.txt command is also invalid because there is no --suid option in chmod. The -v option only means that chmod will output a diagnostic for every file processed. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing Users and Groups; chmod(1) - Linux manual page

NEW QUESTION 68
A junior administrator is setting up a new Linux server that is intended to be used as a router at a remote site. Which of the following parameters will accomplish this goal?
A.

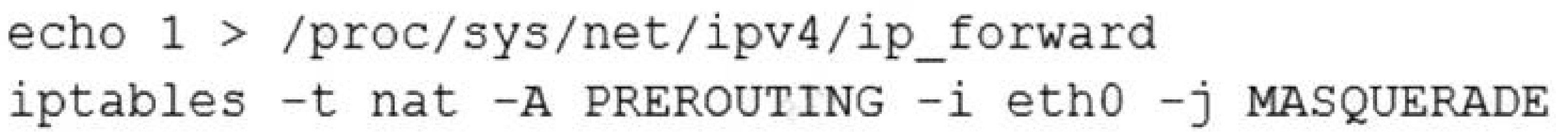

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A PREROUTING -i eth0 -j MASQUERADE
```

A.

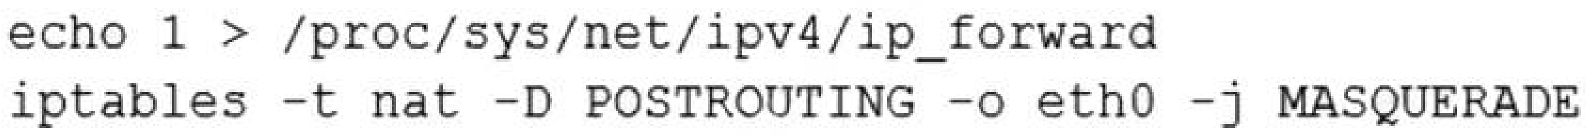

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -D POSTROUTING -o eth0 -j MASQUERADE
```

B.

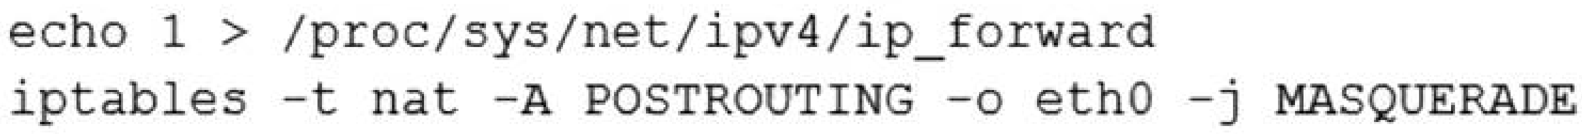

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
```

C.

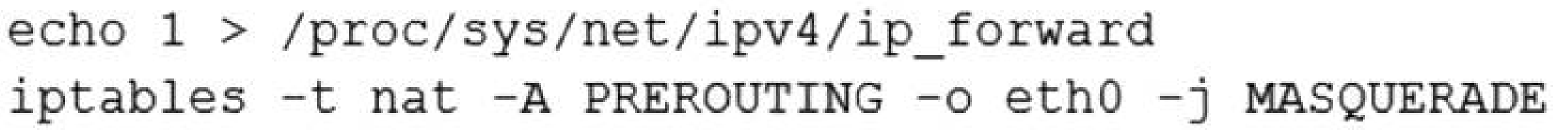

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A PREROUTING -o eth0 -j MASQUERADE
```

Answer: C

Explanation:
The parameter net.ipv4.ip_forward=1 will accomplish the goal of setting up a new Linux server as a router. This parameter enables the IP forwarding feature, which allows the server to forward packets between different network interfaces. This is necessary for a router to route traffic between different networks. The parameter can be set
in the /etc/sysctl.conf file or by using the sysctl command. This is the correct parameter to use to accomplish the goal. The other options are incorrect because they either do not exist (net.ipv4.ip_forwarding or net.ipv4.ip_route) or do not enable IP forwarding (net.ipv4.ip_forward=0). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 382.

NEW QUESTION 71
Which of the following actions are considered good security practices when hardening a Linux server? (Select two).

A. Renaming the root account to something else
B. Removing unnecessary packages
C. Changing the default shell to /bin/csh
D. Disabling public key authentication
E. Disabling the SSH root login possibility
F. Changing the permissions on the root filesystem to 600

Answer: BE

Explanation:
Some good security practices when hardening a Linux server are:
? Removing unnecessary packages (B) to reduce the attack surface and eliminate potential vulnerabilities
? Disabling the SSH root login possibility (E) to prevent unauthorized access and brute-force attacks on the root account References:
? [CompTIA Linux+ Study Guide], Chapter 9: Securing Linux, Section: Hardening Linux
? [How to Harden Your Linux Server]

NEW QUESTION 76
A Linux administrator cloned an existing Linux server and built a new server from that clone. The administrator encountered the following error after booting the cloned server:

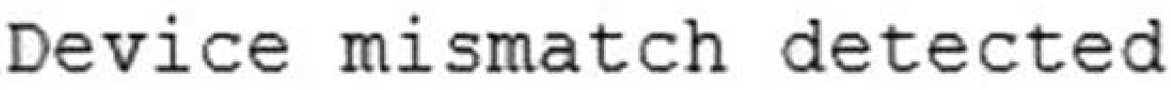

```
Device mismatch detected
```

The administrator performed the commands listed below to further troubleshoot and mount the missing filesystem:

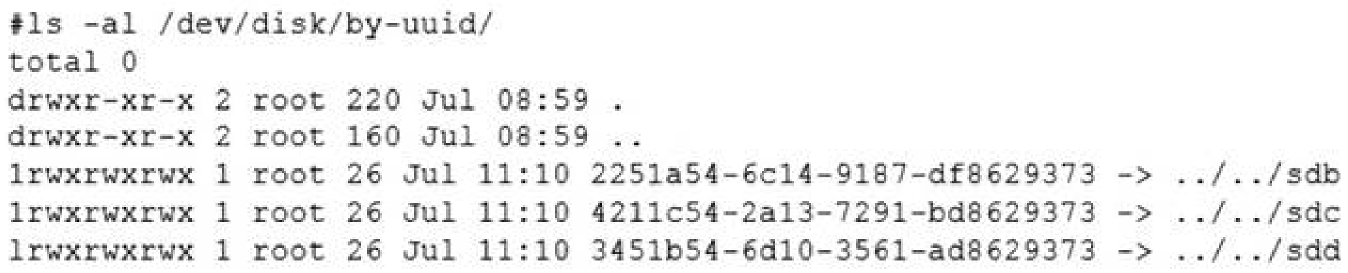

```
#ls -al /dev/disk/by-uuid/
total 0
drwxr-xr-x 2 root 220 Jul 08:59 .
drwxr-xr-x 2 root 160 Jul 08:59 ..
lrwxrwxrwx 1 root 26 Jul 11:10 2251a54-6c14-9187-df8629373 -> ../../sdb
lrwxrwxrwx 1 root 26 Jul 11:10 4211c54-2a13-7291-bd8629373 -> ../../sdc
lrwxrwxrwx 1 root 26 Jul 11:10 3451b54-6d10-3561-ad8629373 -> ../../sdd
```

Which of the following should administrator use to resolve the device mismatch issue and mount the disk?

A. mount disk by device-id
B. fsck -A
C. mount disk by-label
D. mount disk by-blkid

Answer: A

Explanation:
The administrator should use the command mount disk by device-id to resolve the device mismatch issue and mount the disk. The issue is caused by the cloned server having a different device name for the disk than the original server. The output of blkid shows that the disk has the device name /dev/sdb1 on the cloned server, but the output of cat /etc/fstab shows that the disk is expected to have the device name /dev/sda1. The command mount disk by device-id will mount the disk by using its unique identifier (UUID) instead of its device name. The UUID can be obtained from the output of blkid or lsblk -f. The command will mount the disk to the specified mount point (/data) and resolve the issue. The other options are incorrect because they either do not mount the disk (fsck -A), do not use the correct identifier (mount disk by-label or mount disk by-blkid), or do not exist (mount disk by-blkid). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 318-319.

NEW QUESTION 81
An administrator is trying to diagnose a performance issue and is reviewing the following output:

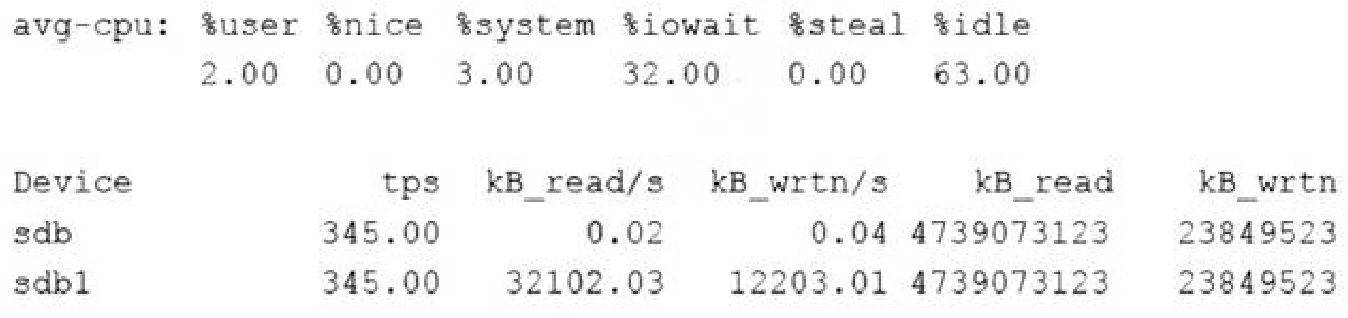

```
avg-cpu: %user %nice %system %iowait %steal %idle
          2.00  0.00   3.00   32.00   0.00  63.00


Device          tps  kB_read/s  kB_wrtn/s    kB_read    kB_wrtn
sdb          345.00       0.02       0.04 4739073123   23849523
sdb1         345.00   32102.03   12203.01 4739073123   23849523
```

System Properties: CPU: 4 vCPU
Memory: 40GB
Disk maximum IOPS: 690
Disk maximum throughput: 44Mbps | 44000Kbps
Based on the above output, which of the following BEST describes the root cause?

A. The system has reached its maximum IOPS, causing the system to be slow.
B. The system has reached its maximum permitted throughput, therefore iowait is increasing.
C. The system is mostly idle, therefore the iowait is high.
D. The system has a partitioned disk, which causes the IOPS to be doubled.

Answer: B

Explanation:
The system has reached its maximum permitted throughput, therefore iowait
is increasing. The output of iostat -x shows that the device sda has an average throughput of 44.01 MB/s, which is equal to the disk maximum throughput of 44 Mbps. The output also shows that the device sda has an average iowait of 99.99%, which means that the CPU is waiting for the disk to complete the I/O requests. This indicates that the disk is the bottleneck and the system is slow due to the high iowait. The other options are incorrect because they are not supported by the outputs. The system has not reached its maximum IOPS, as the device sda has an average IOPS of 563.50, which is lower than the disk maximum IOPS of 690. The system is not mostly idle, as the output of top shows that the CPU is 100% busy. The system does not have a partitioned disk, as the output of lsblk shows that the device sda has only one partition sda1. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Optimizing Linux Systems, pages 513-514.

NEW QUESTION 84
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 86
A Linux administrator is adding a new configuration file to a Git repository. Which of the following describes the correct order of Git commands to accomplish the task successfully?

A. pull -> push -> add -> checkout
B. pull -> add -> commit -> push
C. checkout -> push -> add -> pull
D. pull -> add -> push -> commit

Answer: B

Explanation:
The correct order of Git commands to add a new configuration file to a Git repository is pull -> add -> commit -> push. The pull command will fetch and merge the changes from the remote repository to the local repository, ensuring that the local repository is up to date. The add command will stage the new configuration file for the next commit, marking it as a new file to be tracked by Git. The commit command will create a new snapshot of the project state with the new configuration file and a descriptive message. The push command will publish the commit to the remote repository, updating the remote branch with the new configuration file. The pull -> push -> add -> checkout order is incorrect, as it will not create a commit for the new configuration file, and it will switch to a different branch without pushing the changes. The checkout -> push -> add -> pull order is incorrect, as it will switch to a different branch before adding the new configuration file, and it will overwrite the local changes with the remote changes without creating a commit. The pull -> add -> push -> commit order is incorrect, as it will not create a commit before pushing the changes, and it will create a commit that is not synchronized with the remote branch. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 20: Writing and Executing Bash Shell Scripts, page 579.

NEW QUESTION 90
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 94
A systems administrator is adding a Linux-based server and removing a Windows-based server from a cloud-based environment. The changes need to be validated before they are applied to the cloud-based environment. Which of the following tools should be used to meet this requirement?

A. Ansible
B. git clone
C. git pull
D. terraform plan

Answer: D

Explanation:
Terraform is a tool for building, changing, and managing infrastructure as code in a cloud- based environment. Terraform uses configuration files to describe the desired state of the infrastructure and applies changes accordingly. Terraform supports various cloud providers, such as AWS, Azure, Google Cloud Platform, and more.
To validate changes before they are applied to the cloud-based environment, the administrator can use the terraform plan command. This command will compare the current state of the infrastructure with the desired state defined in the configuration files and show what actions will be performed to achieve the desired state. This command will not make any changes to the infrastructure but only show a plan of changes. The statement D is correct.
The statements A, B, and C are incorrect because they do not validate changes before they are applied to the cloud-based environment. Ansible is another tool for automating infrastructure management, but it does not have a plan command. Git clone and git pull are commands for working with git repositories, which are used for version control of code. References: [How to Use Terraform to Manage Cloud Infrastructure]

NEW QUESTION 96
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:

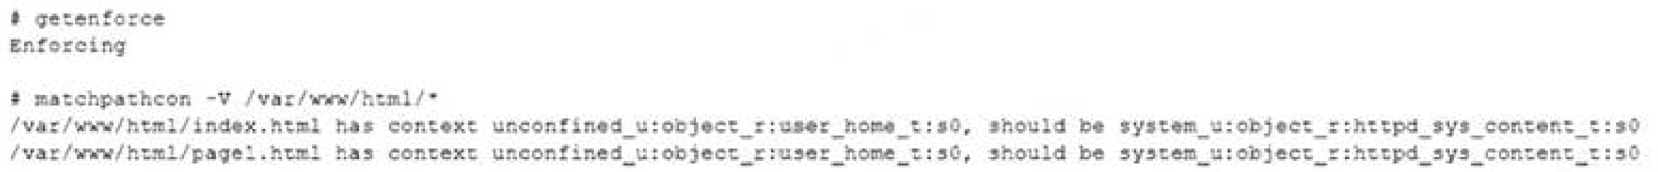

```
# getenforce
Enforcing

# 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 101
A systems administrator is tasked with mounting a USB drive on a system. The USB drive has a single partition, and it has been mapped by the system to the device /dev/sdb. Which of the following commands will mount the USB to /media/usb?

A. mount /dev/sdb1 /media/usb
B. mount /dev/sdb0 /media/usb
C. mount /dev/sdb /media/usb
D. mount -t usb /dev/sdb1 /media/usb

Answer: A

Explanation:
The mount /dev/sdb1 /media/usb command will mount the USB drive to /media/usb. This command will attach the filesystem on the first partition of the USB drive (/dev/sdb1) to the mount point /media/usb, making it accessible to the system. The mount /dev/sdb0 /media/usb command is invalid, as there is no such device as /dev/sdb0. The mount /dev/sdb /media/usb command is incorrect, as it will try to mount the entire USB drive instead of its partition, which may cause errors or data loss. The mount -t usb
/dev/sdb1 /media/usb command is incorrect, as usb is not a valid filesystem type for mount. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Disk Storage, page 455.

NEW QUESTION 103
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 104
DRAG DROP
A new drive was recently added to a Linux system. Using the environment and tokens provided, complete the following tasks:
• Create an appropriate device label.
• Format and create an ext4 file system on the new partition. The current working directory is /.

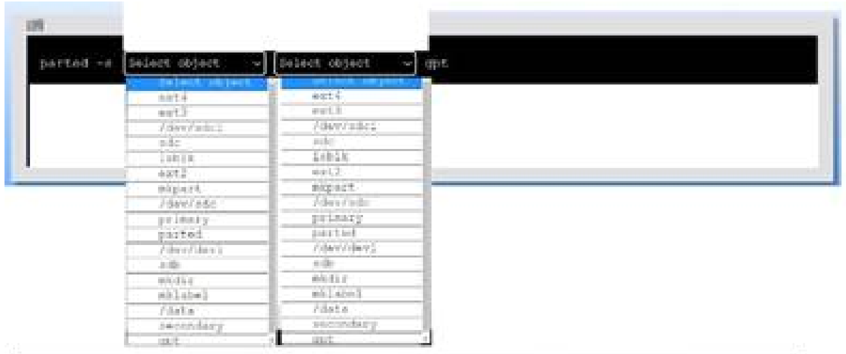

```
parted -s [Select object] [Select object] gpt
```

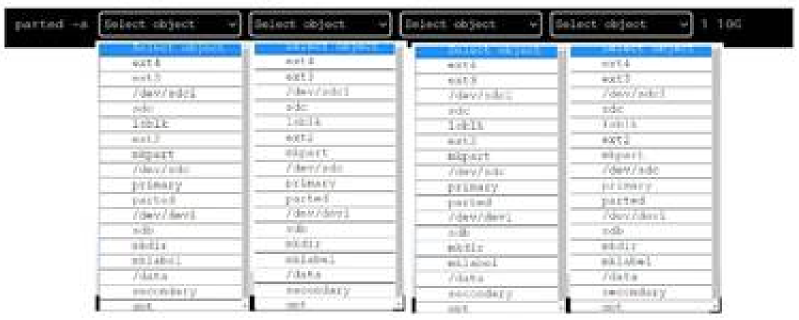

```
parted -s [Select object] [Select object] [Select object] [Select object] 1 10G
```

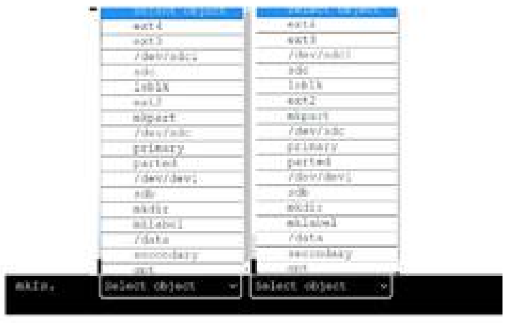

```
mkfs. [Select object] [Select object]
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
To create an appropriate device label, format and create an ext4 file system on the new partition, you can use the following commands:
? To create a GPT (GUID Partition Table) label on the new drive /dev/sdc, you can use the parted command with the -s option (for script mode), the device name (/dev/sdc), the mklabel command, and the label type (gpt). The command is:
parted -s /dev/sdc mklabel gpt
? To create a primary partition of 10 GB on the new drive /dev/sdc, you can use the parted command with the -s option, the device name (/dev/sdc), the mkpart command, the partition type (primary), the file system type (ext4), and the start and end points of the partition (1 and 10G). The command is:
parted -s /dev/sdc mkpart primary ext4 1 10G
? To format and create an ext4 file system on the new partition /dev/sdc1, you can use the mkfs command with the file system type (ext4) and the device name (/dev/sdc1). The command is:
mkfs.ext4 /dev/sdc1
You can verify that the new partition and file system have been created by using the lsblk command, which will list all block devices and their properties.

NEW QUESTION 105
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 110
Users are unable to create new files on the company's FTP server, and an administrator is troubleshooting the issue. The administrator runs the following commands:

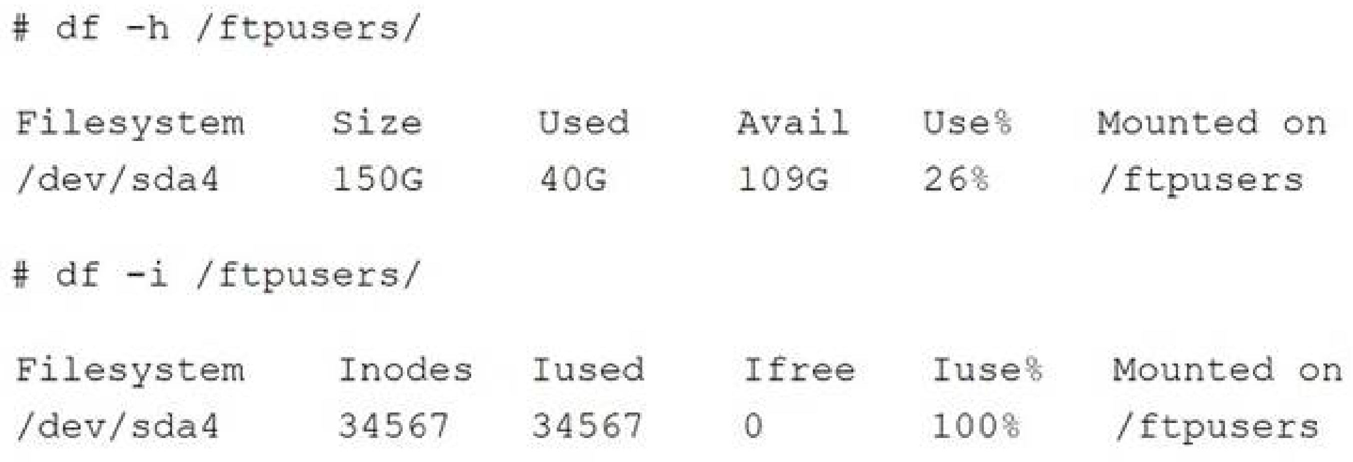

```
# df -h /ftpusers/

Filesystem     Size     Used     Avail    Use%    Mounted on
/dev/sda4      150G     40G      109G     26%     /ftpusers

# df -i /ftpusers/

Filesystem     Inodes   Iused    Ifree    Iuse%   Mounted on
/dev/sda4      34567    34567    0        100%    /ftpusers
```

Which of the following is the cause of the issue based on the output above?

A. The users do not have the correct permissions to create files on the FTP server.
B. The ftpusers filesystem does not have enough space.
C. The inodes is at full capacity and would affect file creation for users.
D. ftpusers is mounted as read only.

Answer: C

Explanation:
The cause of the issue based on the output above is C. The inodes is at full capacity and would affect file creation for users.
An inode is a data structure that stores information about a file or directory, such as its name, size, permissions, owner, timestamps, and location on the disk. Each file or directory has a unique inode number that identifies it. The number of inodes on a filesystem is fixed when the filesystem is created, and it determines how many files and directories can be created on that filesystem. If the inodes are exhausted, no new files or directories can be created, even if there is enough disk space available.
The output for the second command shows that the /ftpusers/ filesystem has 0% of inodes available, which means that all the inodes have been used up. This would prevent users from creating new files on the FTP server. The administrator should either delete some unused files or directories to free up some inodes, or resize the filesystem to increase the number of inodes.
The other options are incorrect because:
* A. The users do not have the correct permissions to create files on the FTP server.
This is not true, because the output for the first command shows that the /ftpusers/ filesystem has 26% of disk space available, which means that there is enough space for users to create files. The permissions of the files and directories are not shown in the output, but they are not relevant to the issue of inode exhaustion.
* B. The ftpusers filesystem does not have enough space.
This is not true, because the output for the first command shows that the /ftpusers/ filesystem has 26% of disk space available, which means that there is enough space for users to create files. The issue is not related to disk space, but to inode capacity.
* D. ftpusers is mounted as read only.
This is not true, because the output for the first command does not show any indication that the /ftpusers/ filesystem is mounted as read only. If it was, it would have an (ro) flag next to the mounted on column. A read only filesystem would prevent users from creating or modifying files on the FTP server, but it would not affect the inode usage.

NEW QUESTION 114
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:

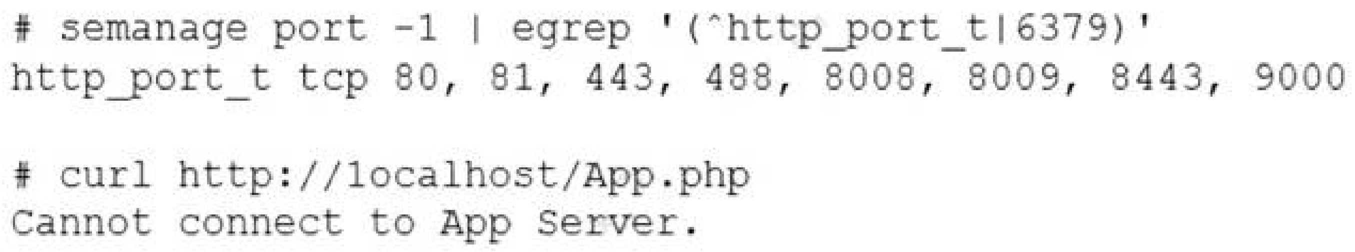

```
# semanage port -l | egrep '(^http_port_t|6379)'
http_port_t tcp 80, 81, 443, 488, 8008, 8009, 8443, 9000

# 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 116
A new Linux systems administrator just generated a pair of SSH keys that should allow connection to the servers. Which of the following commands can be used to copy a key file to remote servers? (Choose two.)

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

Answer: DF

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

NEW QUESTION 121

An administrator needs to increase the system priority of a process with PID 2274. Which of the following commands should the administrator use to accomplish this task?

A. renice —n —15 2274
B. nice -15 2274
C. echo "—15" > /proc/PID/2274/priority
D. ps —ef I grep 2274

Answer: A

Explanation:
The renice command is used to change the priority of a running process by specifying its PID and the new nice value. The -n flag indicates the amount of change in the nice value, which can be positive or negative. A lower nice value means a higher priority, so -15 will increase the priority of the process with PID 2274. The administrator needs to have root privileges to do this.
References:
? The renice command is listed as one of the commands to manipulate process priority in the web search result 1.
? The renice command is also explained with examples in the web search result 2.
? The CompTIA Linux+ Certification Exam Objectives mention that the candidate should be able to “manage process execution priorities” as part of the System Operation and Maintenance domain1.

NEW QUESTION 124
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 126
A systems administrator was tasked with assigning the temporary IP address/netmask 192.168.168.1/255.255.255.255 to the interface eth0 of a Linux server. When adding the address, the following error appears:

```
# ip address add 192.168.168.1/33 dev eth0
Error: any valid prefix is expected rather than "192.168.168.1/33".
```

Based on the command and its output above, which of the following is the cause of the issue?

A. The CIDR value /33 should be /32 instead.
B. There is no route to 192.168.168.1/33.
C. The interface eth0 does not exist.
D. The IP address 192.168.168.1 is already in use.

Answer: A

Explanation:
The cause of the issue is that the CIDR value /33 is invalid for an IPv4 address. The CIDR value represents the number of bits in the network prefix of an IP address, and it can range from 0 to 32 for IPv4 addresses. A CIDR value of /33 would imply a network prefix of more than 32 bits, which is impossible for an IPv4 address. To assign a temporary IP address/netmask of 192.168.168.1/255.255.255.255 to eth0, the CIDR value should be /32 instead, which means a network prefix of 32 bits and a host prefix of 0 bits. There is no route to 192.168.168.1/33 is not the cause of the issue, as the ip address add command does not check the routing table. The interface eth0 does not exist is not the cause of the issue, as the ip address add command would display a different error message if the interface does not exist. The IP address 192.168.168.1 is already in use is not the cause of the issue, as the ip address add command would display a different error message if the IP address is already in use. References: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 13: Networking Fundamentals, page 435.

NEW QUESTION 128
A cloud engineer is asked to copy the file deployment.yaml from a container to the host where the container is running. Which of the following commands can accomplish this task?

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

Answer: B

Explanation:
The command docker cp container_id:/deployment.yaml deployment.yaml can accomplish the task of copying the file deployment.yaml from a container to the host.
The docker command is a tool for managing Docker containers and images. The cp option copies files or directories between a container and the local filesystem. The container_id is the identifier of the container, which can be obtained by using the docker ps command.
The /deployment.yaml is the path of the file in the container, which must be preceded by a slash. The deployment.yaml is the path of the file on the host, which can be relative or absolute. The command docker cp container_id:/deployment.yaml deployment.yaml will copy the file deployment.yaml from the container to the current working directory on the host. This is the correct command to use to accomplish the task. The other options are incorrect because they either use the

wrong syntax (docker cp container_id/deployment.yaml deployment.yaml or docker cp container_id/deployment.yaml local://deployment.yaml) or do not exist (docker cp deployment.yaml local://deployment.yaml). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 567.

NEW QUESTION 132
A Linux administrator needs to connect securely to a remote server in order to install application software. Which of the following commands would allow this connection?

A. scp "ABC-key.pem" root@10.0.0.1
B. sftp rooteiO.0.0.1
C. telnet 10.0.0.1 80
D. ssh -i "ABC-key.pem" root@10.0.0.1
E. sftp "ABC-key.pem" root@10.0.0.1

Answer: D

Explanation:
The command ssh -i "ABC-key.pem" root@10.0.0.1 would allow the administrator to connect securely to the remote server in order to install application software. The ssh command is a tool for establishing secure and encrypted connections between remote systems. The -i option specifies the identity file that contains the private key for key-based authentication. The "ABC-key.pem" is the name of the identity file that contains the private key. The root@10.0.0.1 is the username and the IP address of the remote server. The command ssh -i "ABC-key.pem" root@10.0.0.1 will connect to the remote server using the private key and allow the administrator to install application software. This is the correct command to use to connect securely to the remote server. The other options are incorrect because they either do not use key-based authentication (sftp root@10.0.0.1 or telnet 10.0.0.1 80) or do not use the correct syntax for the command (scp "ABC-key.pem" root@10.0.0.1 instead of scp -i "ABC-key.pem" root@10.0.0.1 or sftp "ABC-key.pem" root@10.0.0.1 instead of sftp -i "ABC-key.pem" root@10.0.0.1). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Implementing Basic Security, page 513.

NEW QUESTION 135
A systems administrator received a request to change a user's credentials. Which of the following commands will grant the request?

A. sudo passwd
B. sudo userde 1
C. sudo chage
D. sudo usermod

Answer: A

Explanation:
This command will allow the systems administrator to change the password of another user account in the system. The sudo prefix will grant the administrator the necessary privileges to perform this action, and the passwd command will prompt for the new password for the specified user. For example, if the administrator wants to change the password of a user named tom, the command will look like this:
sudo passwd tom
The other options are incorrect because:
* B. sudo userdel
This command will delete a user account from the system, not change its credentials. The userdel command removes the user's entry from the /etc/passwd and /etc/shadow files, as well as deletes the user's home directory and mail spool. This is not what the request asked for.
* C. sudo chage
This command will change the password expiration and aging information for a user account, not its credentials. The chage command can be used to set or modify various parameters related to password aging, such as the minimum and maximum number of days between password changes, the number of days before password expiration to issue a warning, and so on. This is not what the request asked for.
* D. sudo usermod
This command will modify various attributes of a user account, such as its login name, home directory, default shell, primary group, and so on. However, it cannot change the user's password directly. To do that, the usermod command requires the -p option followed by an encrypted password string, which is not easy to generate manually. Therefore, this is not a practical way to change a user's credentials.
References:
? How to Change Account Passwords on Linux
? How to Change a Password in Linux for Root and Other Users
? CompTIA Linux+ Certification Exam Objectives

NEW QUESTION 137
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:

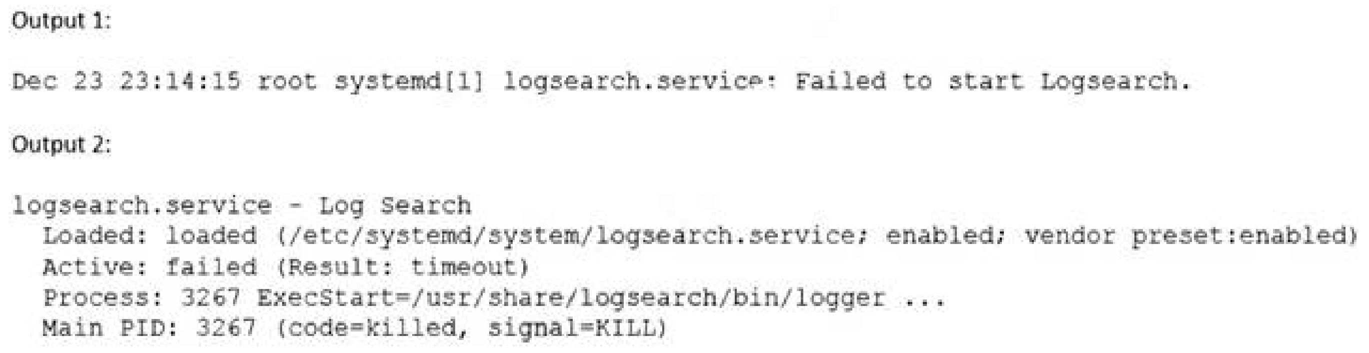

Output 1:

```
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 138
One leg of an LVM-mirrored volume failed due to the underlying physical volume, and a systems administrator is troubleshooting the issue. The following output has been provided:

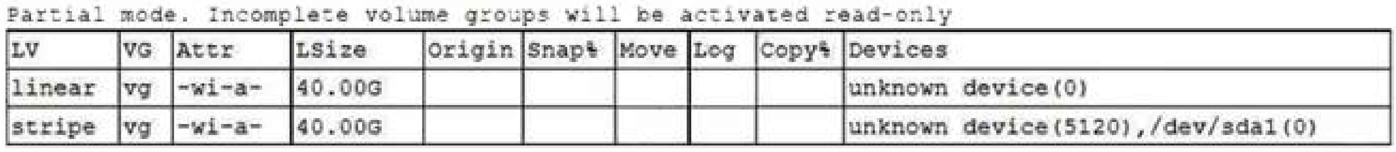

Partial mode. Incomplete volume groups will be activated read-only

LV	VG	Attr	LSize	Origin	Snap%	Move	Log	Copy%	Devices
linear	vg	-wi-a-	40.00G						unknown device(0)
stripe	vg	-wi-a-	40.00G						unknown device(5120),/dev/sda1(0)

Given this scenario, which of the following should the administrator do to recover this volume?

A. Reboot the serve
B. The volume will automatically go back to linear mode.
C. Replace the failed drive and reconfigure the mirror.
D. Reboot the serve
E. The volume will revert to stripe mode.
F. Recreate the logical volume.

Answer: B

Explanation:
The administrator should replace the failed drive and reconfigure the mirror to recover the volume. The LVM (Logical Volume Manager) is a tool for managing disk space on Linux systems. The LVM allows the administrator to create logical volumes that span across multiple physical volumes, such as hard disks or partitions. The LVM also supports different types of logical volumes, such as linear, striped, or mirrored. A mirrored logical volume is a type of logical volume that creates a copy of the data on another physical volume, providing redundancy and fault tolerance. The output shows that the logical volume is mirrored and that one leg of the mirror has failed due to the underlying physical volume. This means that one of the physical volumes that contains the data of the logical volume is damaged or missing. This can cause data loss and performance degradation. The administrator should replace the failed drive and reconfigure the mirror to recover the volume. The administrator should identify the failed physical volume by using commands such as pvdisplay, vgdisplay, or lvdisplay. The administrator should then remove the failed physical volume from the volume group by using the vgreduce command.
The administrator should then install a new drive and create a new physical volume by using the pvcreate command. The administrator should then add the new physical volume to the volume group by using the vgextend command. The administrator should then reconfigure the mirror by using the lvconvert command. The administrator should replace the failed drive and reconfigure the mirror to recover the volume. This is the correct answer to the question. The other options are incorrect because they either do not recover the volume (reboot the server. The volume will automatically go back to linear mode or reboot the server. The volume will revert to stripe mode) or do not preserve the data of the volume (recreate the logical volume). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 333-334.

NEW QUESTION 141
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 145
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 146
An administrator needs to get network information from a group of statically assigned workstations before they are reconnected to the network. Which of the following should the administrator use to obtain this information?

A. ip show
B. ifcfg —a
C. ifcfg —s
D. i fname —s

Answer: B

Explanation:
The ifcfg command is used to configure network interfaces on Linux systems. The -a option displays information about all network interfaces, including their IP addresses, netmasks, gateways, and other parameters. This command can help the administrator obtain the network information from the statically assigned workstations before they are reconnected to the network. References: [Linux Networking: ifcfg Command With Examples]

NEW QUESTION 149
A systems administrator is trying to track down a rogue process that has a TCP listener on a network interface for remote command-and-control instructions. Which of the following commands should the systems administrator use to generate a list of rogue process names? (Select two).

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

Answer: AB

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

NEW QUESTION 152
Users have reported that the interactive sessions were lost on a Linux server. A Linux administrator verifies the server was switched to rescue.target mode for maintenance. Which of the following commands will restore the server to its usual target?

A. telinit 0
B. systemct1 reboot
C. systemct1 get-default
D. systemct1 emergency

Answer: B

Explanation:
The systemct1 reboot command will restore the server to its usual target by rebooting it. This will cause the server to load the default target specified in /etc/systemd/system.conf or /etc/systemd/system/default.target files. The telinit 0 command would shut down the server, not restore it to its usual target. The systemct1 get-default command would display the default target, not change it. The systemct1 emergency command would switch the server to emergency.target mode, which is even more
restrictive than rescue.target mode. References: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 17: System Maintenance and Operation, page 516.

NEW QUESTION 157
A junior systems administrator has just generated public and private authentication keys for passwordless login. Which of the following files will be moved to the

remote servers?

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

Answer: D

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

NEW QUESTION 160
A systems administrator is installing various software packages using a pack-age manager. Which of the following commands would the administrator use on the Linux server to install the package?

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

Answer: D

NEW QUESTION 164
A systems administrator needs to reconfigure a Linux server to allow persistent IPv4 packet forwarding. Which of the following commands is the correct way to accomplish this task?

A. echo 1 > /proc/sys/net/ipv4/ipv_forward
B. sysctl -w net.ipv4.ip_forward=1
C. firewall-cmd --enable ipv4_forwarding
D. systemct1 start ipv4_forwarding

Answer: B

Explanation:
The command sysctl -w net.ipv4.ip_forward=1 enables IPv4 packet forwarding temporarily by setting the kernel parameter net.ipv4.ip_forward to 1. To make this change persistent, the administrator needs to edit the file /etc/sysctl.conf and add the line net.ipv4.ip_forward = 1. The other options are incorrect because they either use the wrong file (/proc/sys/net/ipv4/ipv_forward), the wrong command (firewall- cmd or systemct1), or the wrong option (--enable or start). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 378.

NEW QUESTION 166
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 169
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:

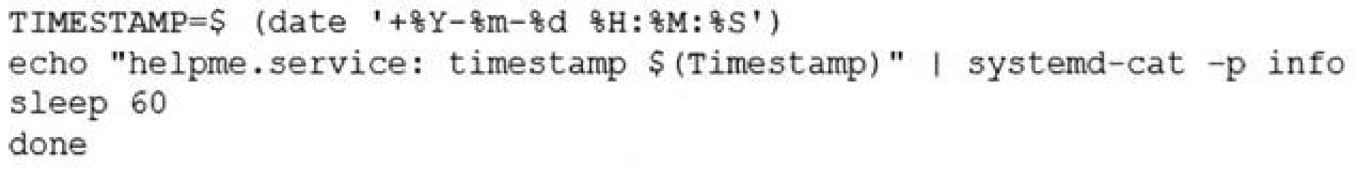

```
TIMESTAMP=$ (date '+%Y-%m-%d %H:%M:%S')
echo "helpme.service: timestamp $(Timestamp)" | systemd-cat -p info
sleep 60
done
```

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 174
A Linux administrator has defined a systemd script docker-repository.mount to mount a volume for use by the Docker service. The administrator wants to ensure that Docker service does not start until the volume is mounted. Which of the following configurations needs to be added to the Docker service definition to best accomplish this task?

A. After=docker-respository.mount
B. ExecStart=/usr/bin/mount -a
C. Requires=docker-repository.mount
D. RequiresMountsFor=docker-repository.mount

Answer: C

Explanation:
This option declares an explicit dependency between the Docker service and the docker- repository.mount unit. It means that the Docker service will not start unless the docker- repository.mount unit is successfully activated. This ensures that the volume is mounted before the Docker service tries to use it12.
References: 1: systemd.unit - systemd unit configuration 2: How to mount host volumes into docker containers in Dockerfile during build

NEW QUESTION 179
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 181
A Linux administrator is trying to start the database service on a Linux server but is not able to run it. The administrator executes a few commands and receives the following output:

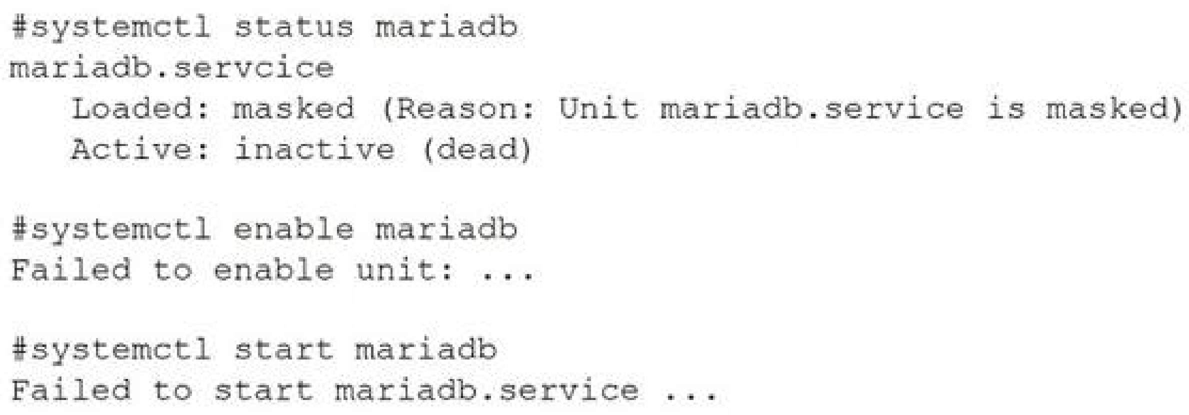

```
#systemctl status mariadb
mariadb.servcice
   Loaded: masked (Reason: Unit mariadb.service is masked)
   Active: inactive (dead)

#systemctl enable mariadb
Failed to enable unit: ...

#systemctl start mariadb
Failed to start mariadb.service ...
```

Which of the following should the administrator run to resolve this issue? (Select two).

A. systemctl unmask mariadb
B. journalctl —g mariadb
C. dnf reinstall mariadb
D. systemctl start mariadb
E. chkconfig mariadb on
F. service mariadb reload

Answer: AD

Explanation:
These commands will unmask the mariadb service, which is currently prevented from starting, and then start it normally. The other commands are either not relevant, not valid, or not sufficient for this task. For more information on how to manage masked services with systemctl, you can refer to the web search result 1.

NEW QUESTION 184
An administrator created an initial Git repository and uploaded the first files. The administrator sees the following when listing the repository:

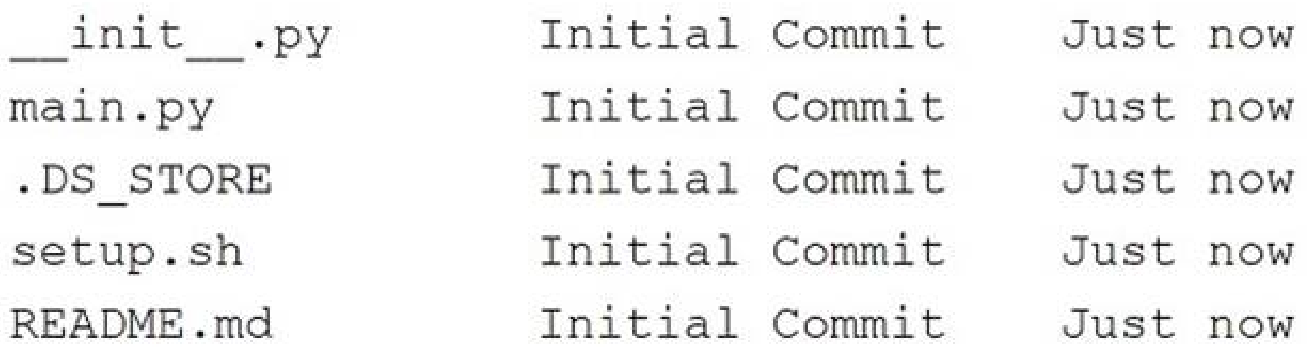

```
__init__.py      Initial Commit    Just now
main.py          Initial Commit    Just now
.DS_STORE        Initial Commit    Just now
setup.sh         Initial Commit    Just now
README.md        Initial Commit    Just now
```

The administrator notices the file . DS STORE should not be included and deletes it from the online repository. Which of the following should the administrator run from the root of the local repository before the next commit to ensure the file is not uploaded again in future commits?

A. rm -f .DS STORE && git push
B. git fetch && git checkout .DS STORE
C. rm -f .DS STORE && git rebase origin main
D. echo .DS STORE >> .gitignore

Answer: D

Explanation:
The correct answer is D. The administrator should run “echo .DS STORE >> .gitignore” from the root of the local repository before the next commit to ensure the file is not uploaded again in future commits.
This command will append the file name .DS STORE to the end of the .gitignore file, which is a special file that tells Git to ignore certain files or directories that should not be tracked or uploaded to the repository. By adding .DS STORE to the .gitignore file, the administrator will prevent Git from staging, committing, or pushing this file in the future.
The other options are incorrect because:
* A. rm -f .DS STORE && git push
This command will delete the file .DS STORE from the local repository and then push the changes to the remote repository. However, this does not prevent the file from being uploaded again in future commits, if it is recreated or copied to the local repository.
* B. git fetch && git checkout .DS STORE
This command will fetch the latest changes from the remote repository and then restore the file .DS STORE from the remote repository to the local repository. This is not what the administrator wants to do, as this will undo the deletion of the file from the online repository.
* C. rm -f .DS STORE && git rebase origin main
This command will delete the file .DS STORE from the local repository and then rebase the local branch onto the main branch of the remote repository. This will rewrite the commit history of the local branch and may cause conflicts or errors. This is not what the administrator wants to do, as this is a risky and unnecessary operation.

NEW QUESTION 185
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 189
Which of the following would significantly help to reduce data loss if more than one drive fails at the same time?

A. Server clustering

B. Load balancing
C. RAID
D. VDI

Answer: C

Explanation:
RAID stands for Redundant Array of Independent Disks, which is a technology that combines multiple physical disks into a logical unit that provides improved performance, reliability, or both. RAID can significantly help to reduce data loss if more than one drive fails at the same time, depending on the RAID level used. For example, RAID 1 (mirroring) duplicates the data on two or more disks, so that if one disk fails, the data can be recovered from another disk. RAID 5 (striping with parity) distributes the data and parity information across three or more disks, so that if one disk fails, the data can be reconstructed from the remaining disks. RAID 6 (striping with double parity) extends RAID 5 by adding another parity block, so that if two disks fail, the data can still be recovered from the remaining disks.
References: [What is RAID?]

NEW QUESTION 193
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 197
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:

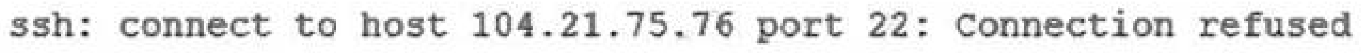

```
ssh: connect to host 104.21.75.76 port 22: Connection refused
```

The administrator reviews the information below:

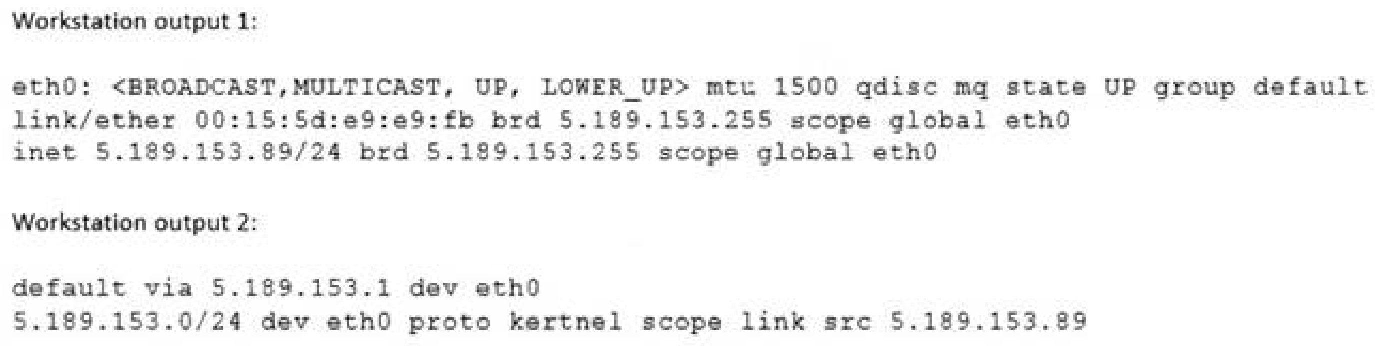

Workstation output 1:

```
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
```

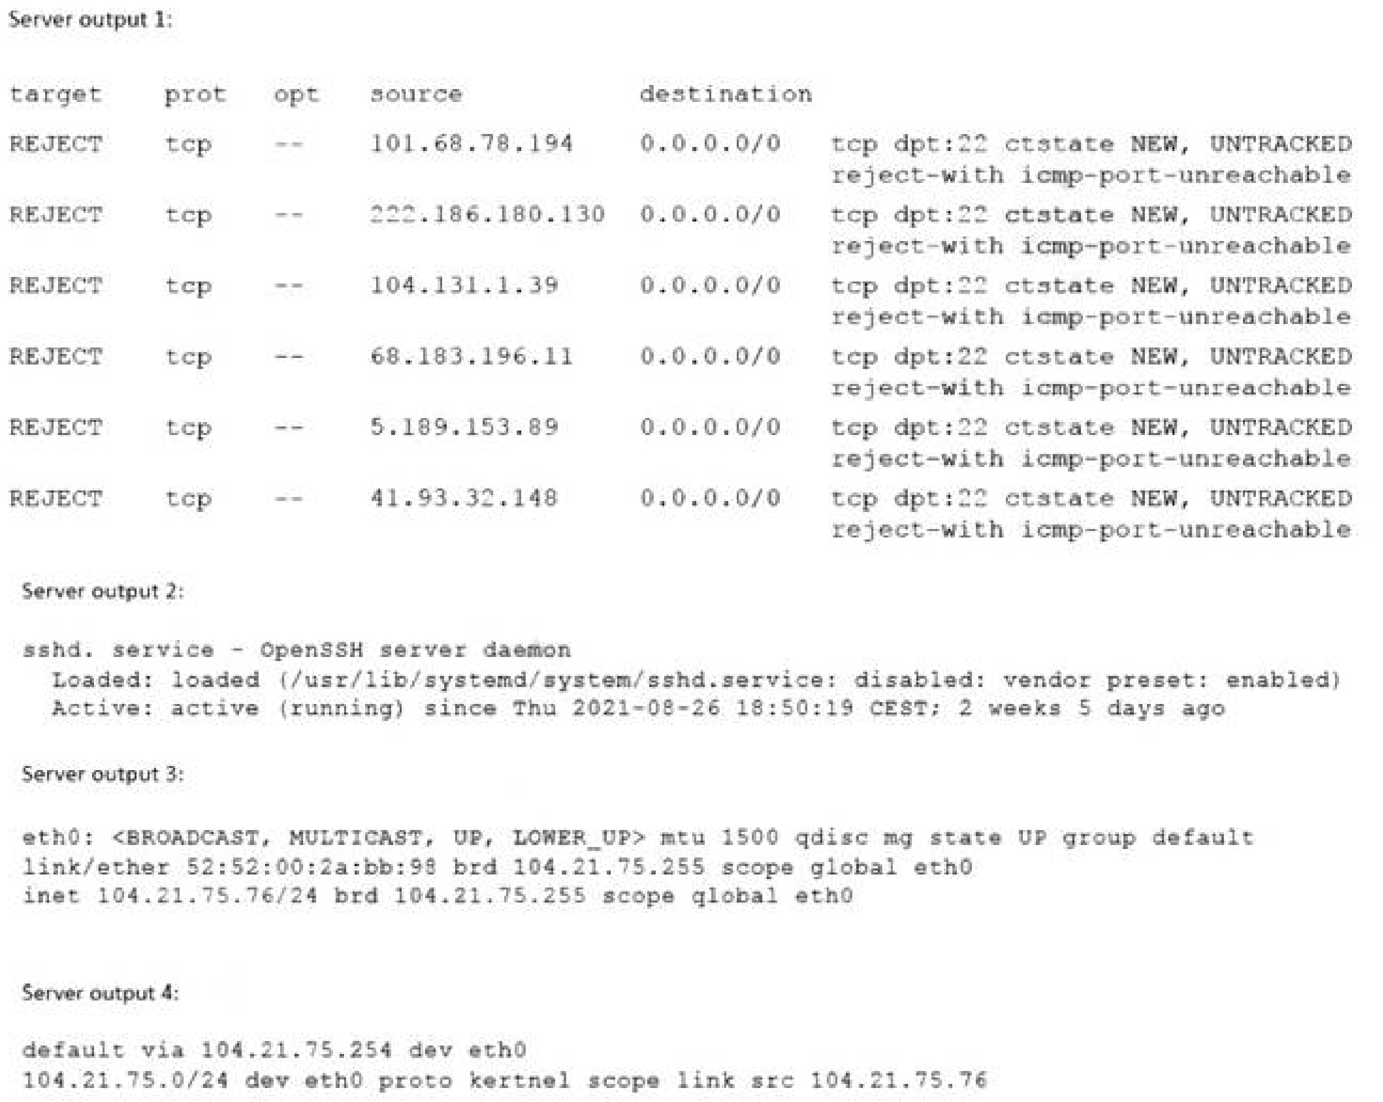

Server output 1:

```
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 202
A Linux administrator is creating a new sudo profile for the accounting user. Which of the following should be added by the administrator to the sudo configuration file so that the accounting user can run / opt/ acc/ report as root?

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

Answer: C

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

NEW QUESTION 204
A user is asking the systems administrator for assistance with writing a script to verify whether a file exists. Given the following:

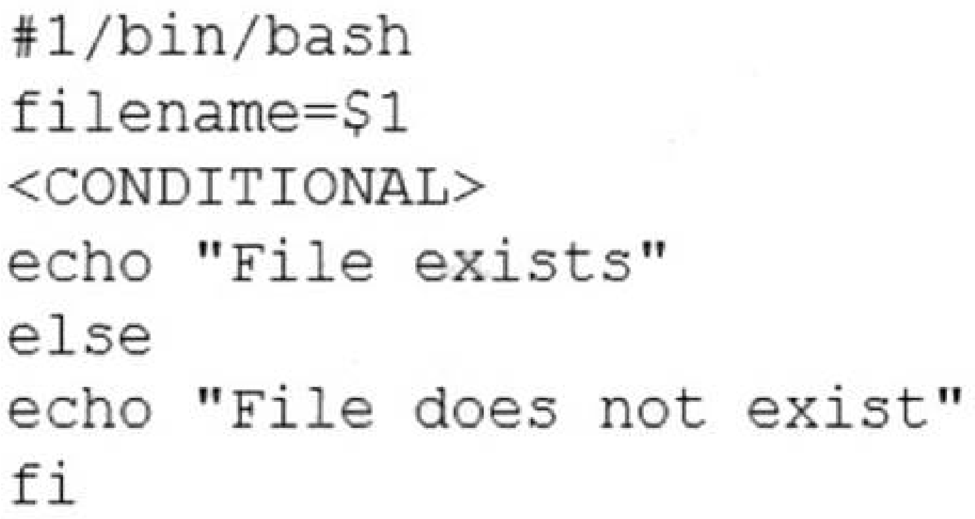

```
#1/bin/bash
filename=$1
<CONDITIONAL>
echo "File exists"
else
echo "File does not exist"
fi
```

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

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

Answer: A

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

NEW QUESTION 205
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 206
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 209
A systems administrator configured firewall rules using firewalld. However, after the system is rebooted, the firewall rules are not present:

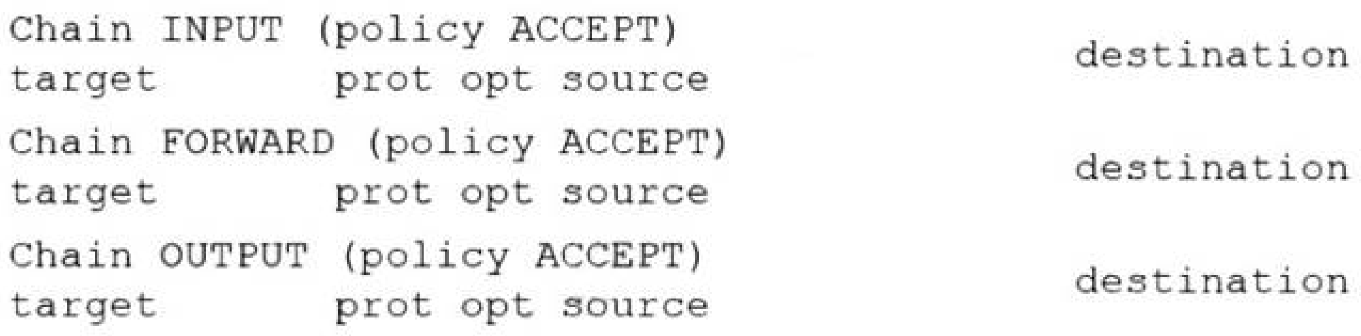

```
Chain INPUT (policy ACCEPT)
target         prot opt source                destination
Chain FORWARD (policy ACCEPT)
target         prot opt source                destination
Chain OUTPUT (policy ACCEPT)
target         prot opt source                destination
```

The systems administrator makes additional checks:

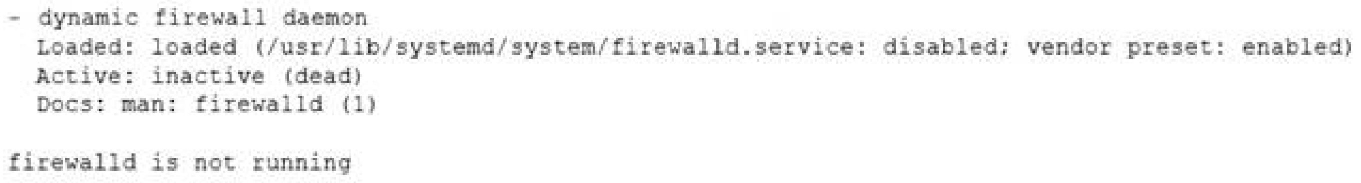

```
- dynamic firewall daemon
  Loaded: loaded (/usr/lib/systemd/system/firewalld.service: disabled; vendor preset: enabled)
  Active: inactive (dead)
  Docs: man: firewalld (1)

firewalld is not running
```

Which of the following is the reason the firewall rules are not active?

A. iptables is conflicting with firewalld.
B. The wrong system target is activated.
C. FIREWALL_ARGS has no value assigned.
D. The firewalld service is not enabled.

Answer: D

Explanation:
The reason the firewall rules are not active is that the firewalld service is not enabled. This means that the service will not start automatically at boot time or after a system reload. To enable the firewalld service, the systems administrator needs to use the command sudo systemct1 enable firewalld. This will create a symbolic link from the firewalld service file to the appropriate systemd target, such as multi-user.target. Enabling the service does not start it immediately, so the systems administrator also needs to use the command sudo systemct1 start firewalld or sudo systemct1 reload firewalld to activate the firewall rules.
The other options are not correct reasons for the firewall rules not being active. iptables is not conflicting with firewalld, because firewalld uses iptables as its backend by default. The wrong system target is not activated, because firewalld is independent of the system target and can be enabled for any target. FIREWALL_ARGS has no value assigned, but this is not a problem, because FIREWALL_ARGS is an optional environment variable that can be used to pass additional arguments to the firewalld daemon, such as --debug or --nofork. If FIREWALL_ARGS is empty or not defined, firewalld will use its default arguments. References: firewalld.service(8) - Linux manual page; firewall-cmd(1) - Linux manual page; systemct1(1) - Linux manual page

NEW QUESTION 210
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 211
The security team has identified a web service that is running with elevated privileges A Linux administrator is working to change the systemd service file to meet security compliance standards. Given the following output:

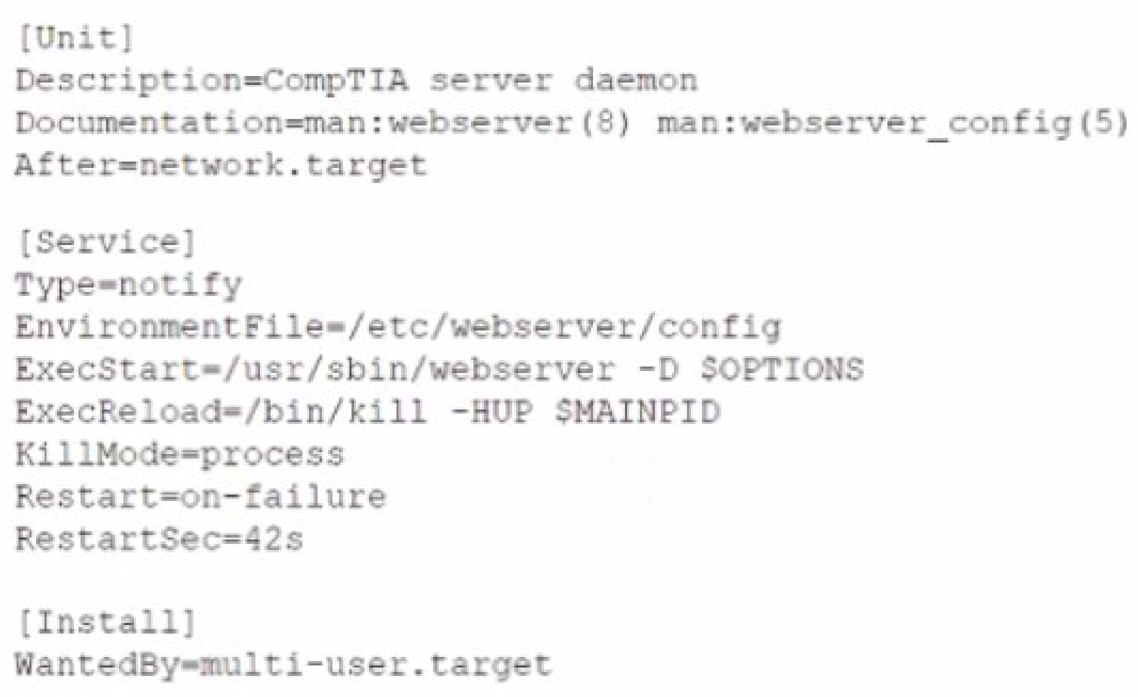

```
[Unit]
Description=CompTIA server daemon
Documentation=man:webserver(8) man:webserver_config(5)
After=network.target

[Service]
Type=notify
EnvironmentFile=/etc/webserver/config
ExecStart=/usr/sbin/webserver -D $OPTIONS
ExecReload=/bin/kill -HUP $MAINPID
KillMode=process
Restart=on-failure
RestartSec=42s

[Install]
WantedBy=multi-user.target
```

Which of the following remediation steps will prevent the web service from running as a privileged user?

A. Removing the ExecStarWusr/sbin/webserver -D SOPTIONS from the service file
B. Updating the Environment File line in the [Service] section to/home/webservice/config
C. Adding the User-webservice to the [Service] section of the service file
D. Changing the:nulti-user.target in the [Install] section to basic.target

Answer: C

Explanation:
The remediation step that will prevent the web service from running as a privileged user is adding the User=webservice to the [Service] section of the service file. The service 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 service file contains various sections and options that specify how the service should be started, stopped, and managed. The [Service] section defines how the service should be executed and what commands should be run. The User option specifies the user name or ID that the service should run as. The webservice is the name of the user that the administrator wants to run the web service as. The administrator should add the User=webservice to the [Service] section of the service file, which will prevent the web service from running as a privileged user, such as root, and improve the security of the system. This is the correct remediation step to use to prevent the web service from running as a privileged user. The other options are incorrect because they either do not change the user that the service runs as (removing the ExecStart=/usr/sbin/webserver -D OPTIONS from the service file or updating the EnvironmentFile line in the [Service] section to /home/webservice/config) or do not affect the user that the service runs as (changing the multi-user.target in the [Install] section to basic.target). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing System Services, page 458.

NEW QUESTION 214
A systems administrator intends to use a UI-JID to mount a new partition per-manently on a Linux system. Which of the following commands can the adminis-trator run to obtain information about the UUlDs of all disks attached to a Linux system?

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

Answer: B

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

NEW QUESTION 218
Several users reported that they were unable to write data to the /oracle1 directory. The following output has been provided:

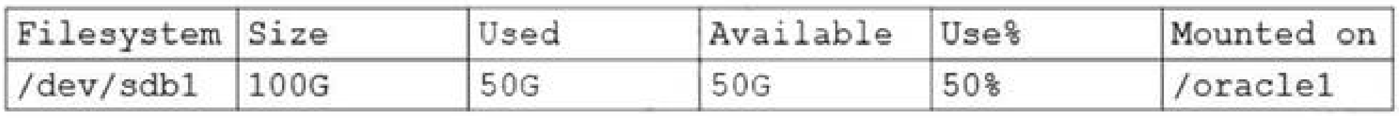

Filesystem	Size	Used	Available	Use%	Mounted on
/dev/sdb1	100G	50G	50G	50%	/oracle1

Which of the following commands should the administrator use to diagnose the issue?

A. df -i /oracle1
B. fdisk -1 /dev/sdb1
C. lsblk /dev/sdb1
D. du -sh /oracle1

Answer: A

Explanation:
The administrator should use the command df -i /oracle1 to diagnose the issue of users being unable to write data to the /oracle1 directory. This command will show the inode usage of the /oracle1 filesystem, which indicates how many files and directories can be created on it. If the inode usage is 100%, it means that no more files or directories can be added, even if there is still free space on the disk. The administrator can then delete some unnecessary files or directories, or increase the inode limit of the filesystem, to resolve the issue.
The other options are not correct commands for diagnosing this issue. The fdisk -l /dev/sdb1 command will show the partition table of /dev/sdb1, which is not

relevant to the inode usage. The lsblk /dev/sdb1 command will show information about /dev/sdb1 as a block device, such as its size, mount point, and type, but not its inode usage. The du -sh /oracle1 command will show the disk usage of /oracle1 in human-readable format, but not its inode usage. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 7: Managing Disk Storage; How to Check Inode Usage in Linux - Fedingo

NEW QUESTION 221
A new disk was presented to a server as /dev/ sdd. The systems administrator needs to check if a partition table is on that disk. Which of the following commands can show this information?

A. Isscsi
B. fdisk
C. blkid
D. partprobe

Answer: B

Explanation:
The command that can be used to check if a partition table is on a disk is fdisk. The fdisk command can display, create, delete, and modify partitions on a disk. To show the partition table of a disk, the administrator can use fdisk -l /dev/sdd (B). References:
? [CompTIA Linux+ Study Guide], Chapter 5: Managing Filesystems and Logical Volumes, Section: Partitioning Disks
? [How to Use Fdisk Command in Linux]

NEW QUESTION 223
......