XK0-005 Dumps

CompTIA Linux+ Certification Exam

https://www.certleader.com/XK0-005-dumps.html

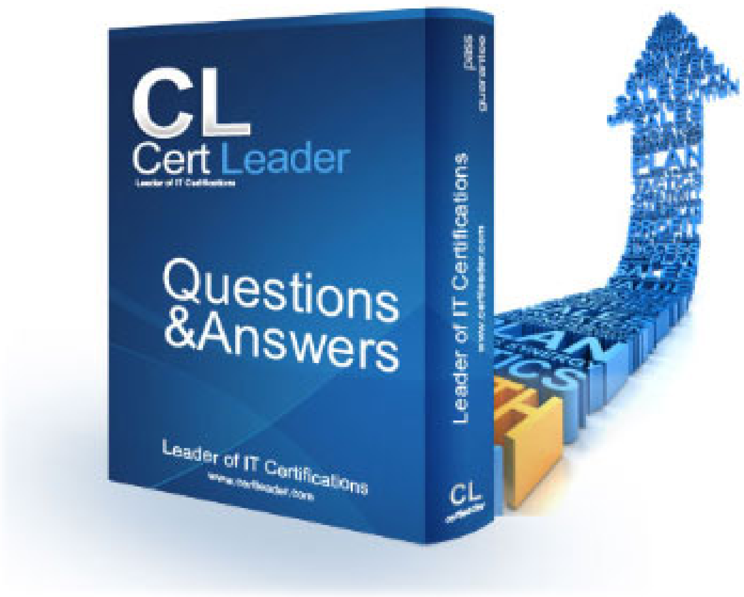

NEW QUESTION 1
A Linux administrator intends to start using KVM on a Linux server. Which of the following commands will allow the administrator to load the KVM module as well as any related dependencies?

A. modprobe kvm
B. insmod kvm
C. depmod kvm
D. hotplug kvm

Answer: A

Explanation:
This command will load the KVM module as well as any related dependencies, such as kvm-intel or kvm-amd, depending on the processor type. The modprobe command is a Linux utility that reads the /etc/modules.conf file and adds or removes modules from the kernel. It also resolves any dependencies between modules, so that they are loaded in the correct order.
The other options are incorrect because:
* B. insmod kvm
This command will only load the KVM module, but not any related dependencies. The insmod command is a low-level Linux utility that inserts a single module into the kernel. It does not resolve any dependencies between modules, so they have to be loaded manually.
* C. depmod kvm
This command will not load the KVM module at all, but only create a list of module dependencies for modprobe to use. The depmod command is a Linux utility that scans the installed modules and generates a file called modules.dep that contains dependency information for each module.
* D. hotplug kvm
This command is invalid and does not exist. The hotplug mechanism is a feature of the Linux kernel that allows devices to be added or removed while the system is running. It does not have anything to do with loading modules.

NEW QUESTION 2
A systems administrator wants to back up the directory /data and all its contents to /backup/data on a remote server named remote. Which of the following commands will achieve the desired effect?

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

Answer: C

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

NEW QUESTION 3
A non-privileged user is attempting to use commands that require elevated account permissions, but the commands are not successful. Which of the following most likely needs to be updated?

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

Answer: C

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

NEW QUESTION 4
A Linux administrator is alerted to a storage capacity issue on a server without a specific mount point or directory. Which of the following commands would be MOST helpful for troubleshooting? (Choose two.)

A. parted

B. df
C. mount
D. du
E. fdisk
F. dd
G. ls

Answer: BD

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

NEW QUESTION 5
A Linux administrator is configuring a new internal web server fleet. The web servers are up and running but can only be reached by users directly via IP address. The administrator is attempting to fix this inconvenience by requesting appropriate records from the DNS team. The details are:
Hostname: devel.comptia.org
IP address: 5.5.5.1, 5.5.5.2, 5.5.5.3, 5.5.5.4
Name server: 5.5.5.254
Additional names: dev.comptia.org, development.comptia.org
Which of the following types of DNS records should the Linux administrator request from the DNS team? (Select three).

A. MX
B. NS
C. PTR
D. A
E. CNAME
F. RRSIG
G. SOA
H. TXT
I. SRV

Answer: BDE

Explanation:
The Linux administrator should request the following types of DNS records from the DNS team:
? A: This record type is used to map a hostname to an IPv4 address. The administrator needs four A records for devel.comptia.org, one for each IP address (5.5.5.1, 5.5.5.2, 5.5.5.3, 5.5.5.4). This will allow users to access the web servers by using the hostname devel.comptia.org instead of the IP addresses1.
? CNAME: This record type is used to create an alias for another hostname. The administrator needs two CNAME records, one for dev.comptia.org and one for development.comptia.org, both pointing to devel.comptia.org. This will allow users to access the web servers by using any of these three hostnames interchangeably1.
? NS: This record type is used to delegate a domain or a subdomain to another name server. The administrator needs one NS record for comptia.org, pointing to 5.5.5.254, which is the name server that hosts the records for the subdomain devel.comptia.org2. This will allow users to resolve the hostnames under comptia.org by querying the name server 5.5.5.2542.
The other record types are not relevant for the administrator’s task:
? MX: This record type is used to specify the mail exchange server for a domain or a subdomain1. The administrator does not need this record type because the web servers are not intended to handle email traffic.
? PTR: This record type is used to map an IP address to a hostname, which is the reverse of an A record1. The administrator does not need this record type because the web servers are not expected to be accessed by their IP addresses.
? RRSIG: This record type is used to provide digital signatures for DNSSEC, which is a security extension for DNS that verifies the authenticity and integrity of DNS responses3. The administrator does not need this record type because it is not mentioned in the task requirements.
? SOA: This record type is used to provide information about the authoritative name server and other parameters for a domain or a subdomain1. The administrator does not need this record type because it is usually created automatically by the name server software when a new zone file is created4.
? TXT: This record type is used to store arbitrary text data that can be used for various purposes, such as SPF, DKIM, DMARC, etc1. The administrator does not need this record type because it is not related to the web server functionality.
? SRV: This record type is used to specify the location and port number of a service that runs on a domain or a subdomain1. The administrator does not need this record type because the web servers use the standard HTTP port 80, which does not require an SRV record.
References: 1: DNS Record Types – CompTIA Network+ N10-007 – 1.8 2: NS Record - DNSimple Help 3: DNSSEC - Wikipedia 4: SOA Record - DNSimple Help

NEW QUESTION 6
An administrator has source code and needs to rebuild a kernel module. Which of the following command sequences is most commonly used to rebuild this type of module?

A. ./configure makemake install
B. wget gcccp
C. tar xvzf buildcp
D. build install configure

Answer: A

Explanation:
The best command sequence to rebuild a kernel module from source code is A. ./configure make make install. This is the standard way to compile and install a Linux kernel module, as explained in the web search result 5. The other commands are either not relevant, not valid, or not sufficient for this task. For example:
? B. wget gcc cp will try to download, compile, and copy a file, but it does not specify the source code, the module name, or the destination directory.
? C. tar xvzf build cp will try to extract, build, and copy a compressed file, but it does not specify the file name, the module name, or the destination directory.
? D. build install configure will try to run three commands that are not defined or recognized by the Linux shell.

NEW QUESTION 7
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:

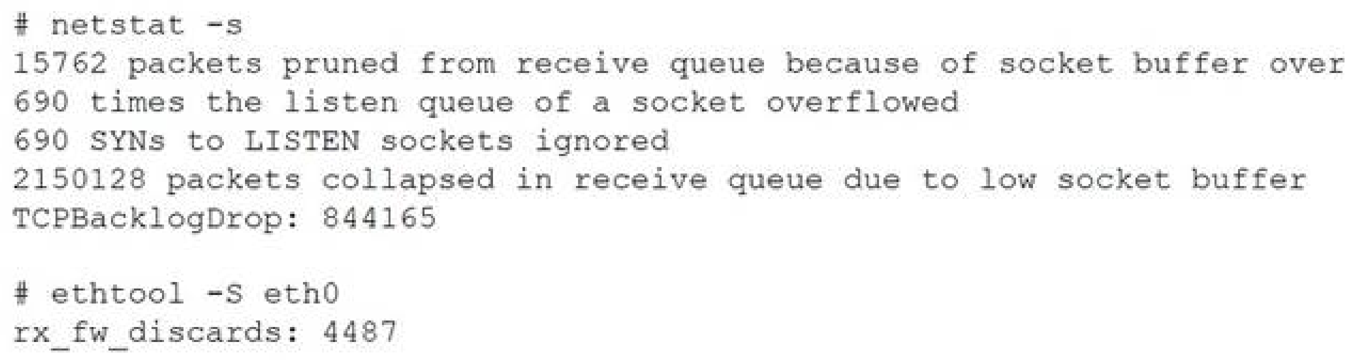

```
# netstat -s
15762 packets pruned from receive queue because of socket buffer over
690 times the listen queue of a socket overflowed
690 SYNs to LISTEN sockets ignored
2150128 packets collapsed in receive queue due to low socket buffer
TCPBacklogDrop: 844165

# 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 8
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:

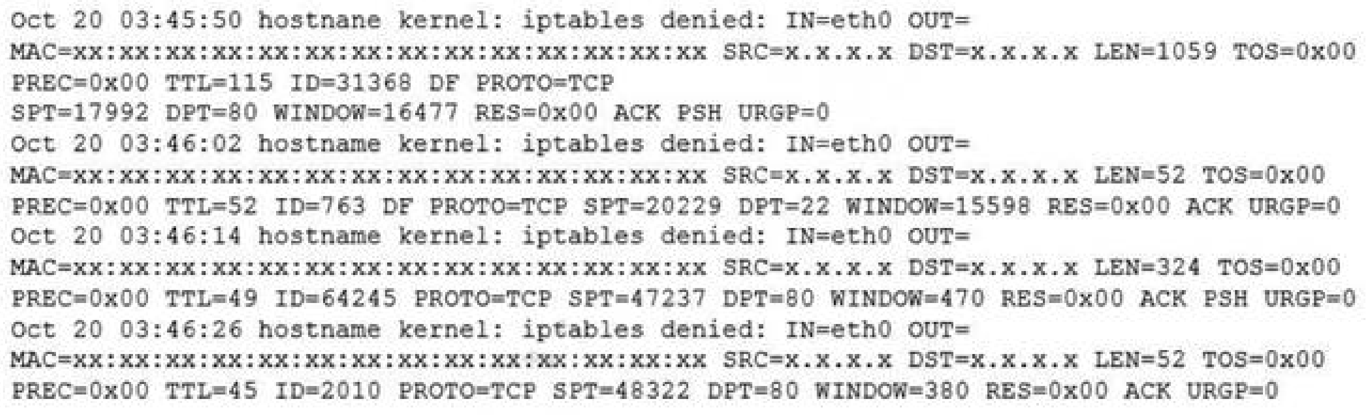

```
Oct 20 03:45:50 hostnane kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx SRC=x.x.x.x DST=x.x.x.x LEN=1059 TOS=0x00
PREC=0x00 TTL=115 ID=31368 DF PROTO=TCP
SPT=17992 DPT=80 WINDOW=16477 RES=0x00 ACK PSH URGP=0
Oct 20 03:46:02 hostname kernel: iptables denied: IN=eth0 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.

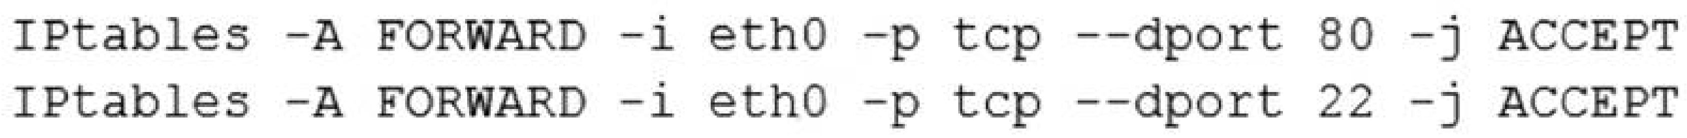

```
IPtables -A FORWARD -i eth0 -p tcp --dport 80 -j ACCEPT
IPtables -A FORWARD -i eth0 -p tcp --dport 22 -j ACCEPT
```

B.

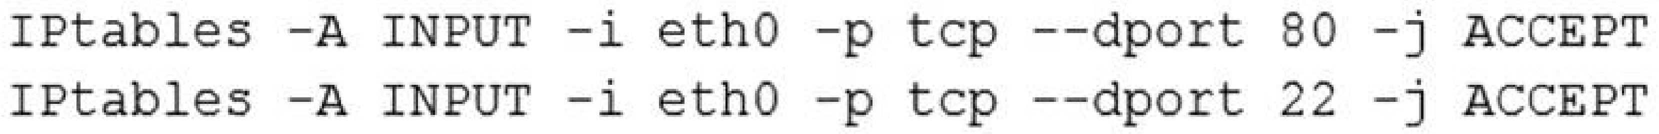

```
IPtables -A INPUT -i eth0 -p tcp --dport 80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --dport 22 -j ACCEPT
```

C.

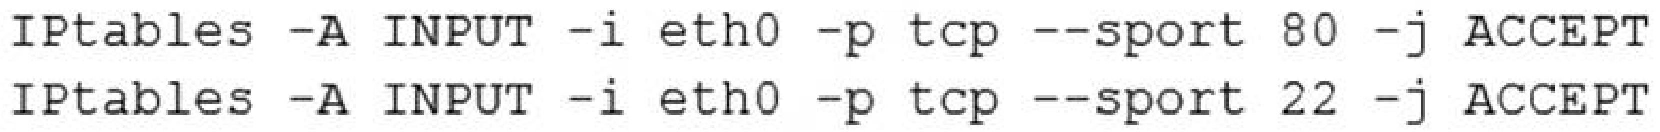

```
IPtables -A INPUT -i eth0 -p tcp --sport 80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --sport 22 -j ACCEPT
```

D.

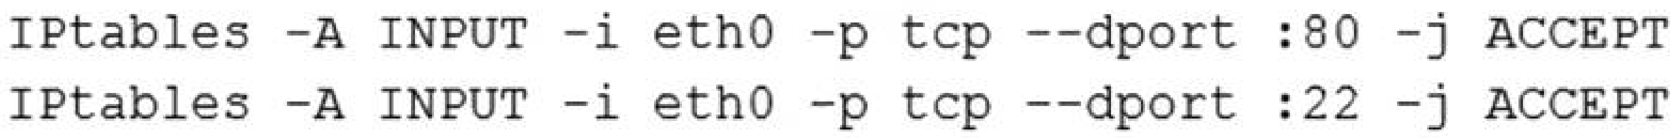

```
IPtables -A INPUT -i eth0 -p tcp --dport :80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --dport :22 -j ACCEPT
```

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 9
An administrator needs to make an application change via a script that must be run only in console mode. Which of the following best represents the sequence the administrator should execute to accomplish this task?

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

Answer: A

Explanation:
The correct answer is A. systemctl isolate multi-user.target sh script.sh systemctl isolate graphical.target
This sequence will allow the administrator to switch from the graphical mode to the console mode, run the script, and then switch back to the graphical mode.
The systemctl command is used to control the systemd system and service manager, which manages the boot targets and services on Linux systems. The isolate subcommand starts the unit specified on the command line and its dependencies and stops all others. The multi-user.target is a boot target that provides a text-based console login, while the graphical.target is a boot target that provides a graphical user interface. By using systemctl isolate, the administrator can change the boot target on the fly without rebooting the system.
The sh command is used to run a shell script, which is a file that contains a series of commands that can be executed by the shell. The script.sh is the name of the script that contains the application change that the administrator needs to make. By running sh script.sh, the administrator can execute the script in the console mode.
The other options are incorrect because:
* B. systemctl isolate graphical.target sh script.sh systemctl isolate multi-user.target
This sequence will switch from the console mode to the graphical mode, run the script, and then switch back to the console mode. This is not what the administrator wants to do, as the script must be run only in console mode.
* C. sh script.sh systemctl isolate multi-user.target systemctl isolate graphical.target
This sequence will run the script in the current mode, which may or may not be console mode, and then switch to console mode and back to graphical mode. This is not what the administrator wants to do, as the script must be run only in console mode.
* D. systemctl isolate multi-user.target systemctl isolate graphical.target sh script.sh
This sequence will switch from graphical mode to console mode and then back to graphical mode, without running the script at all. This is not what the administrator wants to do, as the script must be run only in console mode.
References:
? systemctl(1) - Linux manual page
? How to switch between the CLI and GUI on a Linux server
? How to PROPERLY boot into single user mode in RHEL/CentOS 7/8
? Changing Systemd Boot Target in Linux
? Exit Desktop to Terminal in Ubuntu 19.10

NEW QUESTION 10
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 10
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 14
Users are reporting that writes on a system configured with SSD drives have been taking longer than expected, but reads do not seem to be affected. A Linux systems administrator is investigating this issue and working on a solution. Which of the following should the administrator do to help solve the issue?

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

Answer: A

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

NEW QUESTION 19
A user is unable to remotely log on to a server using the server name server1 and port 22.
The Linux engineer troubleshoots the issue and gathers the following information: Which of the following is most likely causing the issue?

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

Answer: B

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

NEW QUESTION 21
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 26
The group owner of the / home/ test directory would like to preserve all group permissions on files created in the directory. Which of the following commands should the group owner execute?

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

Answer: A

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

symbolic links and operate on all files and directories recursively. The test /home/test arguments are used to specify the new owner and group and the target directory.
References:
? How to Set File Permissions Using chmod
? How to Use Chmod Command in Linux with Examples
? How to Use Chown Command in Linux with Examples
? [How to Use Chgrp Command in Linux with Examples]

NEW QUESTION 30
A developer has been unable to remove a particular data folder that a team no longer uses. The developer escalated the issue to the systems administrator. The following output was received:

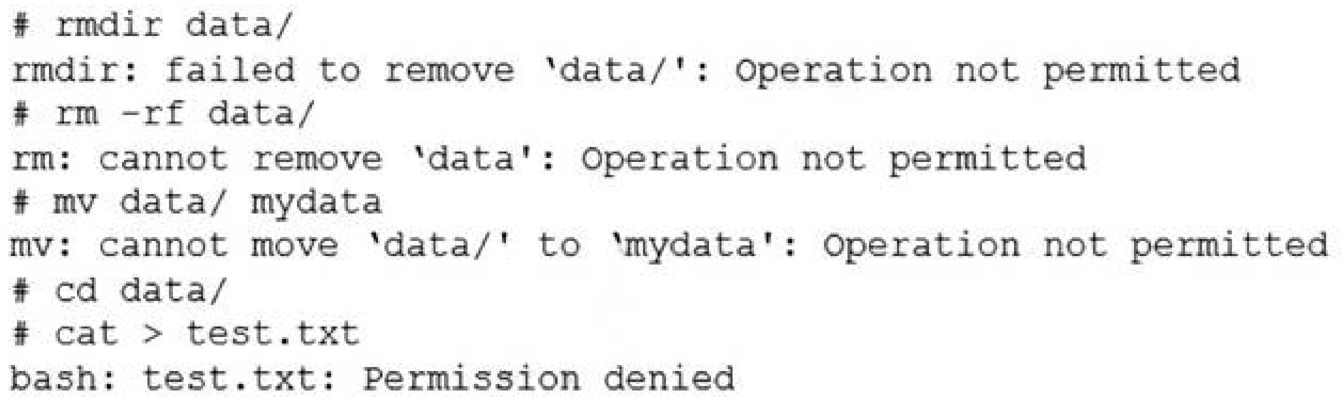

```
# rmdir data/
rmdir: failed to remove 'data/': Operation not permitted
# rm -rf data/
rm: cannot remove 'data': Operation not permitted
# mv data/ mydata
mv: cannot move 'data/' to 'mydata': Operation not permitted
# cd data/
# cat > test.txt
bash: test.txt: Permission denied
```

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

A. chgrp -R 755 data/
B. chmod -R 777 data/
C. chattr -R -i data/
D. chown -R data/

Answer: C

Explanation:
The command that can be used to resolve the issue of being unable to remove a particular data folder is chattr -R -i data/. This command will use the chattr utility to change file attributes on a Linux file system. The -R option means that chattr will recursively change attributes of directories and their contents. The -i option means that chattr will remove (unset) the immutable attribute from files or directories. When a file or directory has the immutable attribute set, it cannot be modified, deleted, or renamed.
The other options are not correct commands for resolving this issue. The chgrp -R 755 data/ command will change the group ownership of data/ and its contents recursively to 755, which is not a valid group name. The chgrp command is used to change group ownership of files or directories. The chmod -R 777 data/ command will change the file mode bits of data/ and its contents recursively to 777, which means that everyone can read, write, and execute them. However, this will not remove the immutable attribute, which prevents deletion or modification regardless of permissions. The chmod command is used to change file mode bits of files or directories. The chown -R data/ command is incomplete and will produce an error. The chown command is used to change the user and/or group ownership of files or directories, but it requires at least one argument besides the file name. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 7: Managing Disk Storage; chattr(1) - Linux manual page; chgrp(1) - Linux manual page; chmod(1) - Linux manual page; chown(1) - Linux manual page

NEW QUESTION 34
Users have been unable to save documents to /home/tmp/temp and have been receiving the following error:
Path not found
A junior technician checks the locations and sees that /home/tmp/tempa was accidentally created instead of /home/tmp/temp. Which of the following commands should the technician use to fix this issue?

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

Answer: B

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

NEW QUESTION 36
A Linux administrator is trying to remove the ACL from the file /home/user/data. txt but receives the following error message:

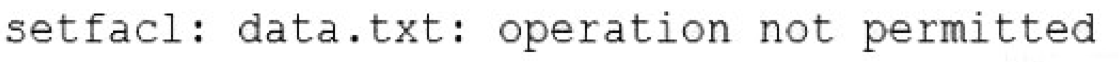

```
setfacl: data.txt: operation not permitted
```

Given the following analysis:

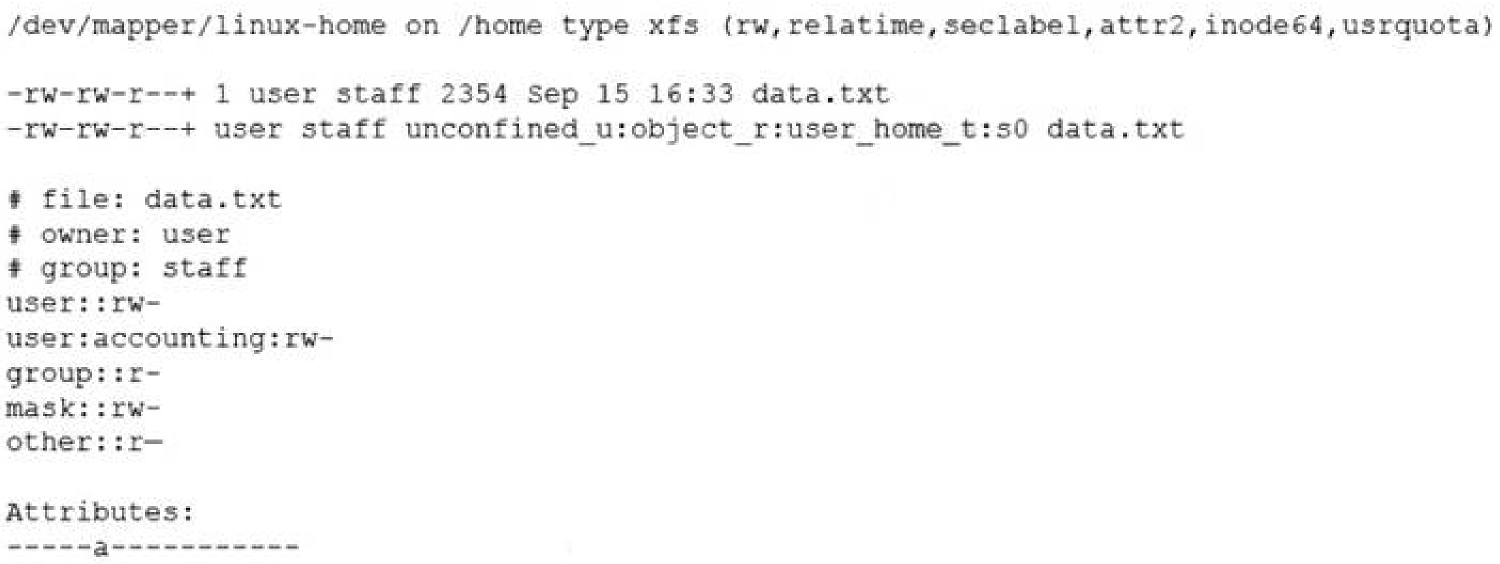

```
/dev/mapper/linux-home on /home type xfs (rw,relatime,seclabel,attr2,inode64,usrquota)

-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 40
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.

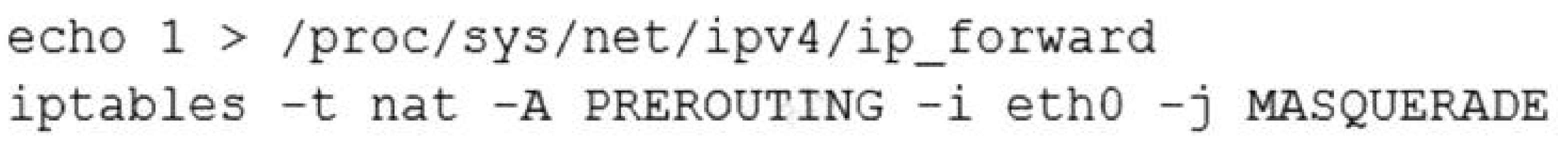

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A PREROUTING -i eth0 -j MASQUERADE
```

A.

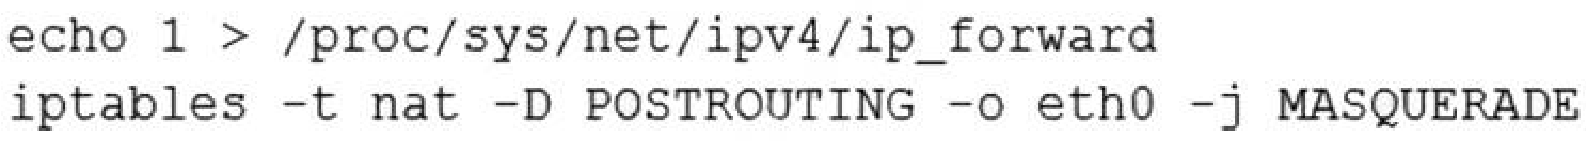

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -D POSTROUTING -o eth0 -j MASQUERADE
```

B.

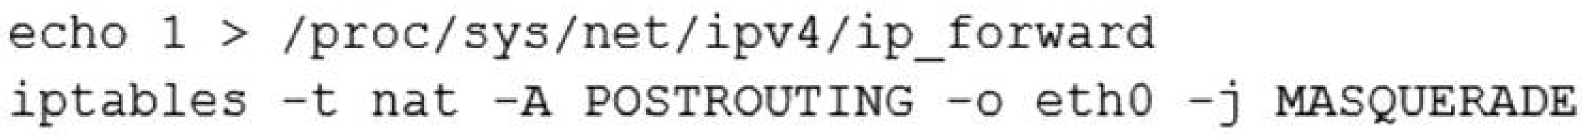

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
```

C.

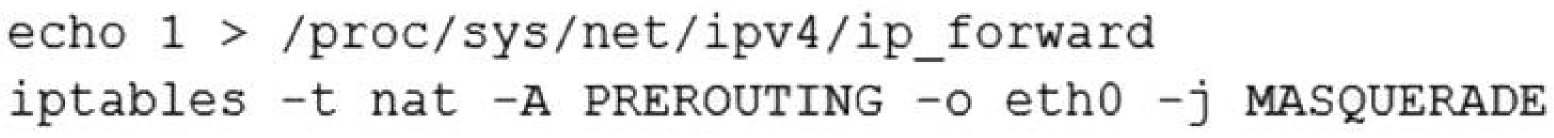

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A PREROUTING -o eth0 -j MASQUERADE
```

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 41
A systems administrator is tasked with preventing logins from accounts other than root, while the file /etc/nologin exists. Which of the following PAM modules will accomplish this task?

A. pam_login.so
B. pam_access.so
C. pam_logindef.so
D. pam_nologin.so

Answer: D

Explanation:
The PAM module pam_nologin.so will prevent logins from accounts other than root, while the file /etc/nologin exists. This module checks for the existence of the file /etc/nologin and displays its contents to the user before denying access. The root user is exempt from this check and can still log in. This is the correct module to accomplish the task. The other options are incorrect because they are either non-existent modules (pam_login.so or pam_logindef.so) or do not perform the required function (pam_access.so controls access based on host, user, or time). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Users and Groups, page 471.

NEW QUESTION 42
An administrator installed an application from source into /opt/operations1/ and has received numerous reports that users are not able to access the application without having to use the full path /opt/operations1/bin/*. Which of the following commands should be used to resolve this issue?

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

Answer: A

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

NEW QUESTION 45
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 48
An administrator is trying to diagnose a performance issue and is reviewing the following output:

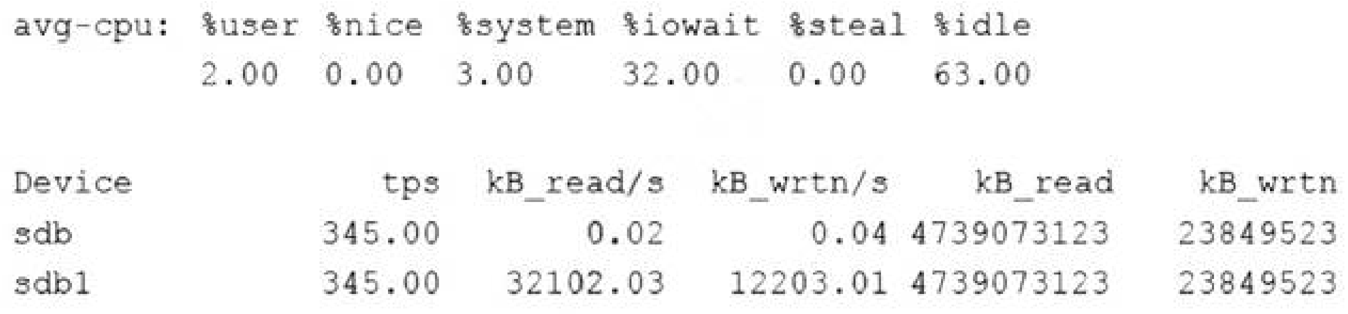

```
avg-cpu: %user %nice %system %iowait %steal %idle
          2.00  0.00   3.00    32.00   0.00  63.00


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 51
A systems administrator is tasked with installing GRUB on the legacy MBR of the SATA hard drive. Which of the following commands will help the administrator accomplish this task?

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

Answer: B

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

NEW QUESTION 54
A systems administrator is compiling a report containing information about processes that are listening on the network ports of a Linux server. Which of the following commands will allow the administrator to obtain the needed information?

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

Answer: A

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

NEW QUESTION 55
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 /.

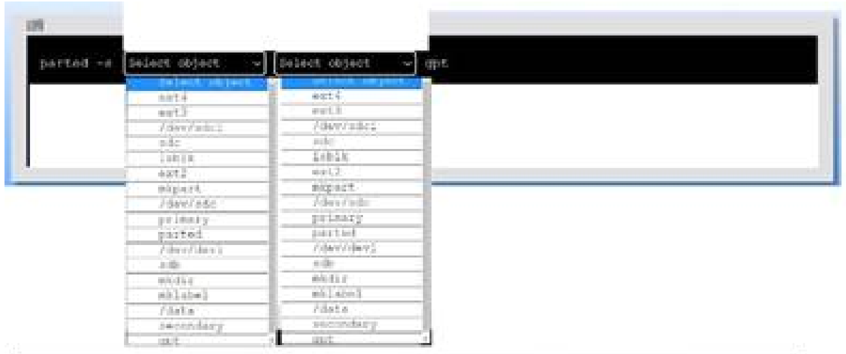

parted -s [Select object] [Select object] gpt

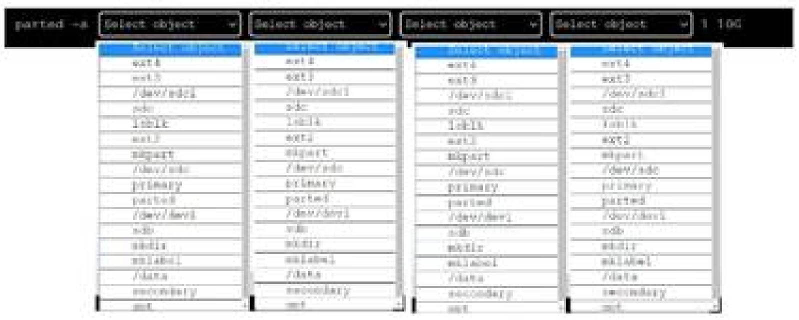

parted -s [Select object] [Select object] [Select object] [Select object] 1 10G

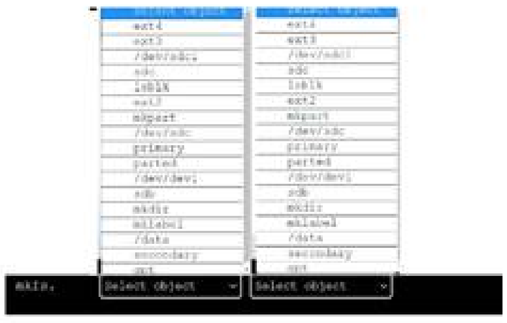

mkfs. [Select object] [Select object]

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 58
Which of the following will prevent non-root SSH access to a Linux server?

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

Answer: A

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

NEW QUESTION 63
What is the main objective when using Application Control?

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

Answer: D

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

NEW QUESTION 64
A systems administrator is encountering performance issues. The administrator runs 3 commands with the following output

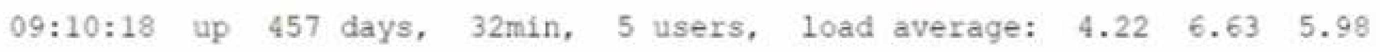

```
09:10:18  up  457 days,  32min,  5 users,  load average:  4.22  6.63  5.98
```

The Linux server has the following system properties CPU: 4 vCPU
Memory: 50GB
Which of the following accurately describes this situation?

A. The system is under CPU pressure and will require additional vCPUs
B. The system has been running for over a year and requires a reboot.
C. Too many users are currently logged in to the system
D. The system requires more memory

Answer: A

Explanation:
Based on the output of the image sent by the user, the system is under CPU pressure and will require additional vCPUs. The output shows that there are four processes running upload.sh scripts that are consuming a high percentage of CPU time (99.7%, 99.6%, 99.5%, and 99.4%). The output also shows that the system has only 4 vCPUs, which means that each process is using almost one entire vCPU. This indicates that the system is struggling to handle the CPU load and may experience performance issues or slowdowns. Adding more vCPUs to the system would help to alleviate the CPU pressure and improve the system performance. The system has not been running for over a year, as the uptime command shows that it has been up for only 1 day, 2 hours, and 13 minutes. The number of users logged in to the system is not relevant to the performance issue, as they are not consuming significant CPU resources. The system does not require more memory, as the free command shows that it has plenty of available memory (49 GB total, 48 GB free). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Memory and Process Execution, pages 468-469.

NEW QUESTION 66
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 68
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 71
A Linux administrator is creating a primary partition on the replacement hard drive for an application server. Which of the following commands should the administrator issue to verify the device name of this partition?

A. sudo fdisk /dev/sda
B. sudo fdisk -s /dev/sda
C. sudo fdisk -l
D. sudo fdisk -h

Answer: C

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

Study Guide, Chapter 10: Managing Storage, page 317.

NEW QUESTION 75
A systems administrator creates a public key for authentication. Which of the following tools is most suitable to use when uploading the key to the remote servers?

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

Answer: B

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

NEW QUESTION 80
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:

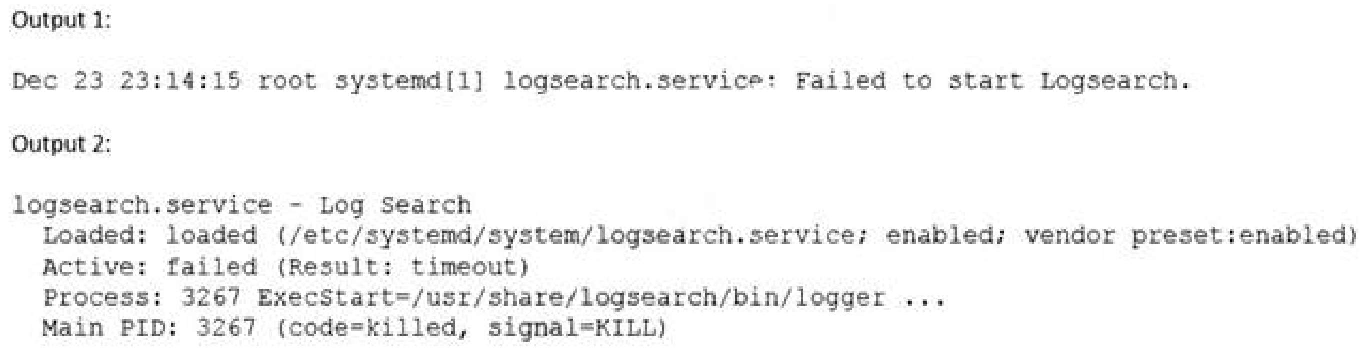

Output 1:

```
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 84
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:

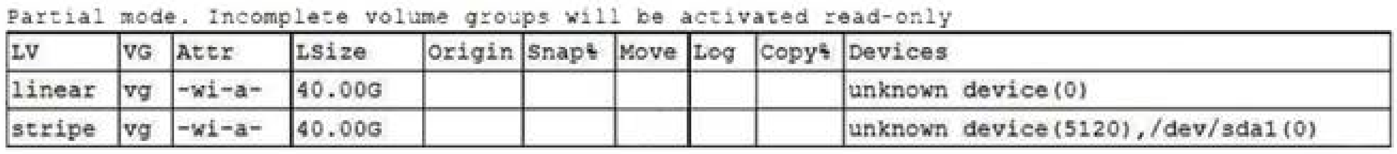

Partial mode. Incomplete volume groups will be activated read-only

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 85
A Linux engineer needs to download a ZIP file and wants to set the nice of value to -10 for this new process. Which of the following commands will help to accomplish the task?

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

Answer: D

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

NEW QUESTION 88
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 91
Which of the following technologies provides load balancing, encryption, and observability in containerized environments?

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

Answer: D

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

NEW QUESTION 93
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 98
A systems administrator is tasked with creating a cloud-based server with a public IP address.

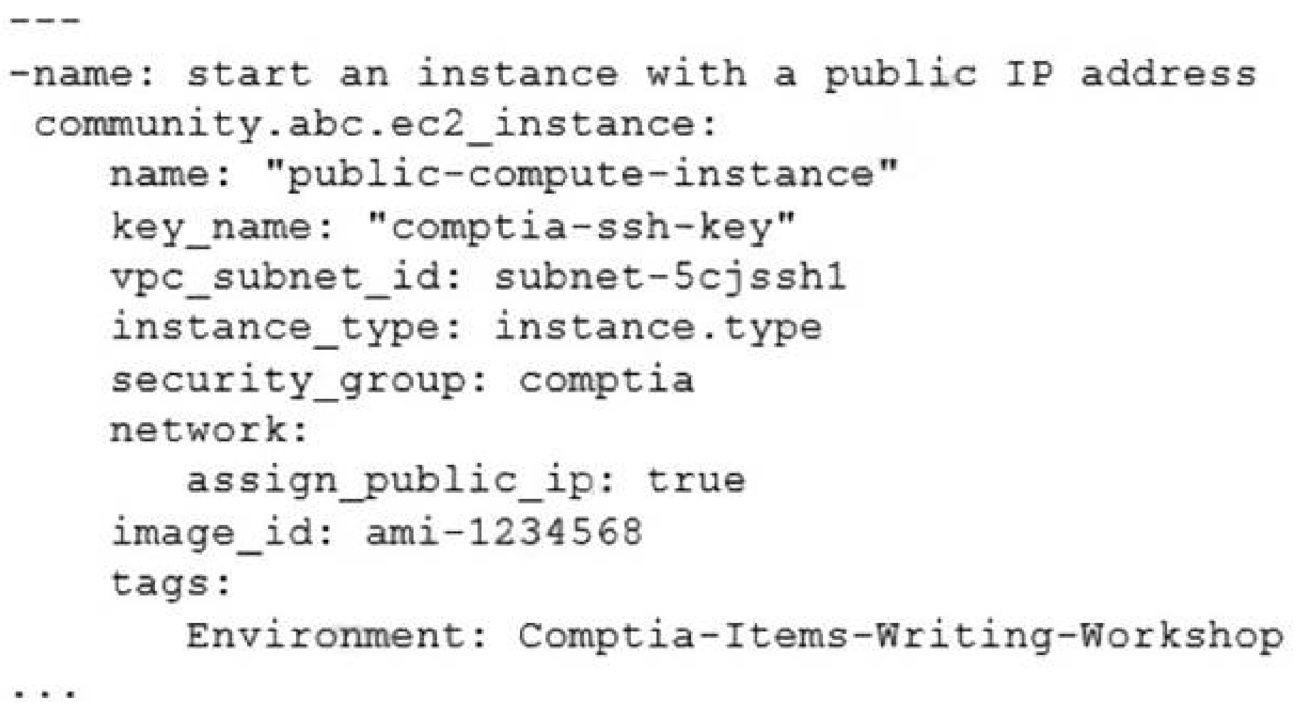

```
---
-name: start an instance with a public IP address
 community.abc.ec2_instance:
    name: "public-compute-instance"
    key_name: "comptia-ssh-key"
    vpc_subnet_id: subnet-5cjssh1
    instance_type: instance.type
    security_group: comptia
    network:
       assign_public_ip: true
    image_id: ami-1234568
    tags:
       Environment: Comptia-Items-Writing-Workshop
...
```

Which of the following technologies did the systems administrator use to complete this task?

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

Answer: D

Explanation:
The systems administrator used Terraform to create a cloud-based server with a public IP address. Terraform is a tool for building, changing, and versioning infrastructure as code. Terraform can create and manage resources on different cloud platforms, such as AWS, Azure, or Google Cloud. Terraform uses a declarative syntax to describe the desired state of the infrastructure and applies the changes accordingly. Terraform can also assign a public IP address to a cloud server by using the appropriate resource attributes. This is the correct technology that the systems administrator used to complete the task. The other options are incorrect because they are either not designed for creating cloud servers (Puppet or Git) or not capable of assigning public IP addresses (Ansible). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 559.

NEW QUESTION 99
An administrator started a long-running process in the foreground that needs to continue without interruption. Which of the following keystrokes should the administrator use to continue running the process in the background?

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

Answer: A

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

NEW QUESTION 104
A DevOps engineer is working on a local copy of a Git repository. The engineer would like to switch from the main branch to the staging branch but notices the

staging branch does not exist. Which of the following Git commands should the engineer use to perform this task?

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

Answer: D

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

NEW QUESTION 108
An administrator added the port 2222 for the SSH server on myhost and restarted the SSH server. The administrator noticed issues during the startup of the service. Given the following outputs:

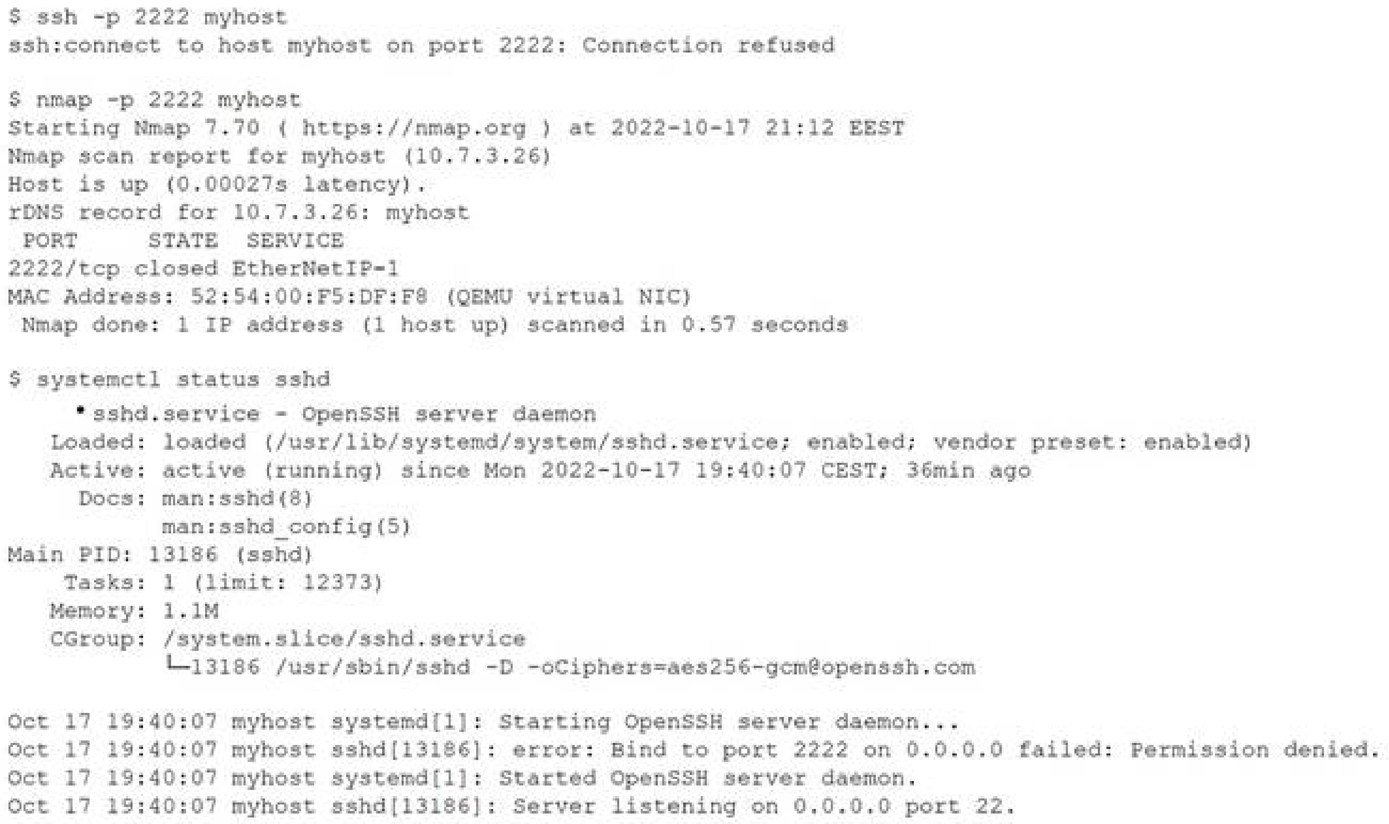

```
$ ssh -p 2222 myhost
ssh:connect to host myhost on port 2222: Connection refused

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

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

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

Which of the following commands will fix the issue?

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

Answer: A

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

issue, as SELinux will still block the SSH server from binding to that port. Moreover, firewalld may not be installed or enabled on some Linux distributions, such as Ubuntu or Debian, which use iptables instead.
References:
? How to configure SSH to use a non-standard port with SELinux set to enforcing
? Change SSH Port on CentOS/RHEL/Fedora With SELinux Enforcing
? How to change SSH port when SELinux policy is enabled

NEW QUESTION 110
A systems administrator is checking the system logs. The administrator wants to look at the last 20 lines of a log. Which of the following will execute the command?

A. tail -v 20
B. tail -n 20
C. tail -c 20
D. tail -l 20

Answer: B

Explanation:
The command tail -n 20 will display the last 20 lines of a file. The -n option specifies the number of lines to show. This is the correct command to execute the task. The other options are incorrect because they either use the wrong options (-v, -c, or -l) or have the wrong arguments (20 instead of 20 filename). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, page 352.

NEW QUESTION 114
Joe, a user, is unable to log in to the Linux system Given the following output:

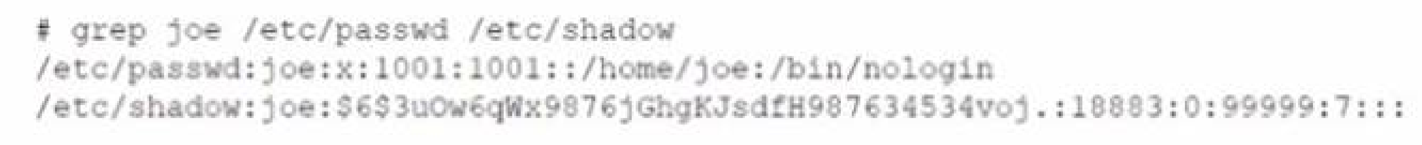

```
# grep joe /etc/passwd /etc/shadow
/etc/passwd:joe:x:1001:1001::/home/joe:/bin/nologin
/etc/shadow:joe:$6$3uOw6qWx9876jGhgKJsdfH987634534voj.:18883:0:99999:7:::
```

Which of the following command would resolve the issue?

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

Answer: B

Explanation:
Based on the output of the image sent by the user, Joe is unable to log in to the Linux system because his account has been locked due to too many failed login attempts. The pam_tally2 -u joe -r command will resolve this issue by resetting Joe’s failed login counter to zero and unlocking his account. This command uses the pam_tally2 module to manage user account locking based on login failures. The usermod -s /bin/bash joe command will change Joe’s login shell to /bin/bash, but this will not unlock his account. The passwd -u joe command will unlock Joe’s password if it has been locked by passwd -l joe, but this will not reset his failed login counter or unlock his account if it has been locked by pam_tally2. The chage -E 90 joe command will set Joe’s account expiration date to 90 days from today, but this will not unlock his account or reset his failed login counter. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 537.

NEW QUESTION 116
A Linux administrator needs to resolve a service that has failed to start. The administrator runs the following command:

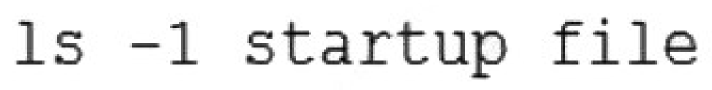

```
ls -1 startup file
```

The following output is returned

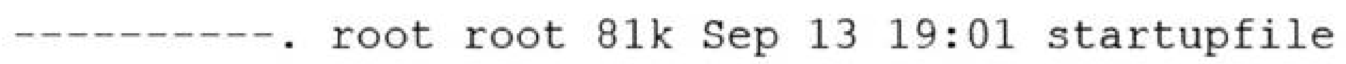

```
----------. root root 81k Sep 13 19:01 startupfile
```

Which of the following is MOST likely the issue?

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

Answer: A

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

NEW QUESTION 117
Several users reported that they were unable to write data to the /oracle1 directory. The following output has been provided:

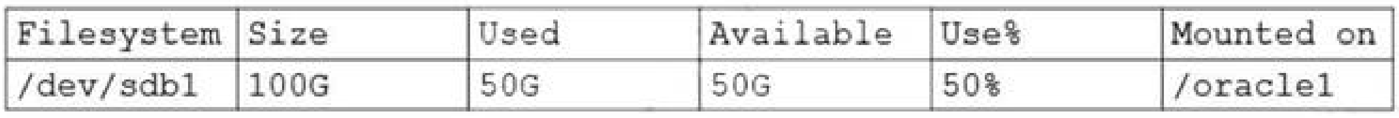

Filesystem	Size	Used	Available	Use%	Mounted on
/dev/sdb1	100G	50G	50G	50%	/oracle1

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 122
A cloud engineer needs to remove all dangling images and delete all the images that do not have an associated container. Which of the following commands will help to accomplish this task?

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

Answer: A

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

NEW QUESTION 124
A Linux administrator is scheduling a system job that runs a script to check available disk space every hour. The Linux administrator does not want users to be able to start the job. Given the following:

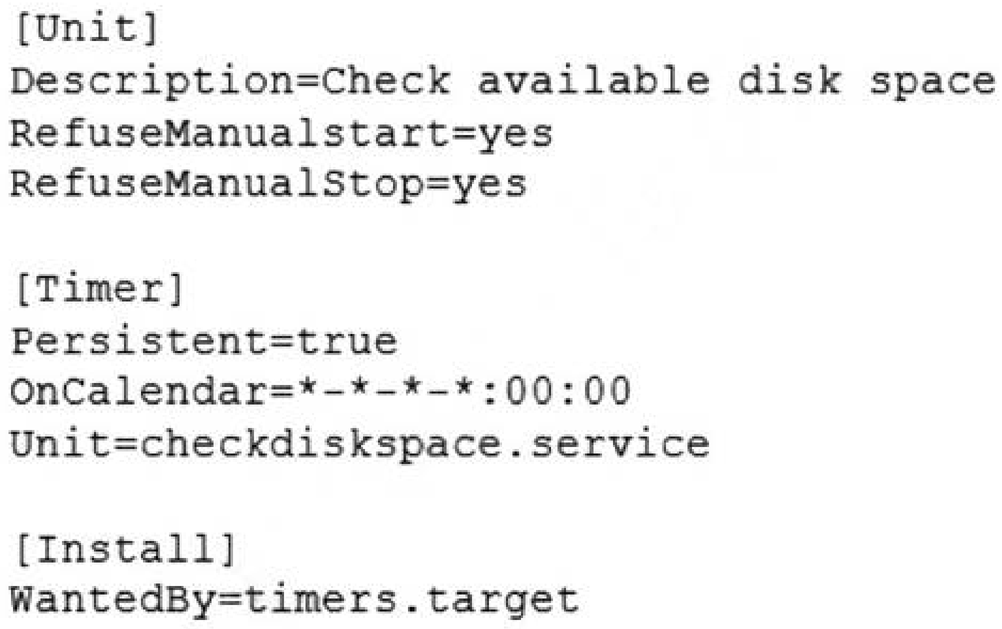

```
[Unit]
Description=Check available disk space
RefuseManualstart=yes
RefuseManualStop=yes

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

[Install]
WantedBy=timers.target
```

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

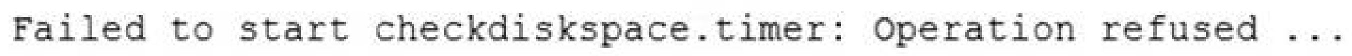

```
Failed to start checkdiskspace.timer: Operation refused ...
```

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

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

Answer: C

Explanation:

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

NEW QUESTION 129
A Linux administrator needs to correct the permissions of a log file on the server. Which of the following commands should be used to set filename.log permissions to -rwxr—r--. ?

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

Answer: A

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

NEW QUESTION 134
Which of the following will prevent non-root SSH access to a Linux server?

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

Answer: A

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

NEW QUESTION 135
Joe, a user, is unable to log in to the Linux system. Given the following output:

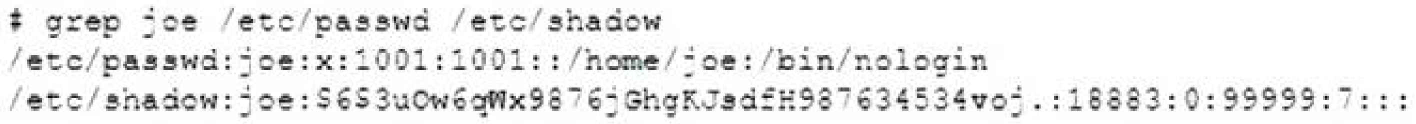

```
# grep joe /etc/passwd /etc/shadow
/etc/passwd:joe:x:1001:1001::/home/joe:/bin/nologin
/etc/shadow:joe:$6$3uOw6qWx9876jGhgKJsdfH987634534voj.:18883:0:99999:7:::
```

Which of the following commands would resolve the issue?

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

Answer: B

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

joe). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Implementing Basic Security, page 517.

NEW QUESTION 137
A Linux administrator booted up the server and was presented with a non-GUI terminal. The administrator ran the command systemct1 isolate graphical.target and rebooted the system by running systemct1 reboot, which fixed the issue. However, the next day the administrator was presented again with a non-GUI terminal. Which of the following is the issue?

A. The administrator did not reboot the server properly.
B. The administrator did not set the default target to basic.target.
C. The administrator did not set the default target to graphical.target.
D. The administrator did not shut down the server properly.

Answer: C

Explanation:
The issue is that the administrator did not set the default target to graphical.target. A target is a unit of systemd that groups together other units by a common purpose or state. The graphical.target is a target that starts the graphical user interface (GUI) along with other services. The administrator used the command systemct1 isolate graphical.target to switch to this target temporarily, but this does not change the default target that is activated at boot time. To make this change permanent, the administrator should have used the command systemct1 set-default graphical.target, which creates a symbolic link from /etc/systemd/system/default.target to /usr/lib/systemd/system/graphical.target.
The other options are not correct explanations for the issue. The administrator did reboot the server properly by using systemct1 reboot, which shuts down and restarts the system cleanly. The administrator did not need to set the default target to basic.target, which is a minimal target that only starts essential services. The administrator did not shut down the server improperly, which could have caused file system corruption or data loss, but not affect the default target. References: systemct1(1) - Linux manual page; How to Change Runlevels (targets) in SystemD

NEW QUESTION 142
A Linux administrator is providing a new Nginx image from the registry to local cache. Which of the following commands would allow this to happen?

A. docker pull nginx
B. docker attach nginx
C. docker commit nginx
D. docker import nginx

Answer: A

Explanation:
The command that would allow this to happen is docker pull nginx. Docker is a software platform that allows the administrator to create, run, and manage containers on Linux systems. Containers are isolated and lightweight environments that can run applications and services without affecting the host system. Docker uses images to create containers, which are files that contain the code, libraries, dependencies, and configuration of the applications and services. Docker uses a registry to store and distribute images, which is a service that hosts and serves images. Docker Hub is the default public registry that provides a large number of official and community images. Nginx is a popular web server and reverse proxy that can run as a container. The command docker pull nginx will download the latest version of the Nginx image from the Docker Hub registry to the local cache, which is the storage location for the images on the host system. This will allow the administrator to provide a new Nginx image from the registry to the local cache. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not download an image from the registry (docker attach nginx or docker commit nginx) or do not exist (docker import nginx). References: CompTIA Linux+ (XK0- 005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 571.

NEW QUESTION 143
A Linux administrator needs to analyze a failing application that is running inside a container. Which of the following commands allows the Linux administrator to enter the running container and analyze the logs that are stored inside?

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

Answer: B

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

NEW QUESTION 147
A Linux systems administrator receives a notification that one of the server's filesystems is full. Which of the following commands would help the administrator to identify this filesystem?

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

Answer: C

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

NEW QUESTION 151
A Linux administrator needs to create a new user named user02. However, user02 must be in a different home directory, which is under /comptia/projects. Which of the following commands will accomplish this task?

A. useradd -d /comptia/projects user02
B. useradd -m /comptia/projects user02
C. useradd -b /comptia/projects user02
D. useradd -s /comptia/projects user02

Answer: A

Explanation:
The command useradd -d /comptia/projects user02 will accomplish the task of creating a new user named user02 with a different home directory.
The useradd command is a tool for creating new user accounts on Linux systems. The - d option specifies the home directory for the new user, which is the directory where the user’s personal files and settings are stored. The /comptia/projects is the path of the home directory for the new user, which is different from the default location of /home/user02.
The user02 is the name of the new user. The command useradd -d /comptia/projects user02 will create a new user named user02 with a home directory under /comptia/projects. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not specify the home directory for the new user (useradd -m /comptia/projects user02 or useradd -s /comptia/projects user02) or do not use the correct option for the home directory (useradd -b /comptia/projects user02 instead of useradd -d /comptia/projects user02). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Users and Groups, page 403.

NEW QUESTION 156
Ann, a security administrator, is performing home directory audits on a Linux server. Ann issues the su Joe command and then issues the Is command. The output displays files that reside in Ann's home directory instead of Joe's. Which of the following represents the command Ann should have issued in order to list Joe's files?

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

Answer: A

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

NEW QUESTION 159
A systems administrator made some unapproved changes prior to leaving the company. The newly hired administrator has been tasked with revealing the system to a compliant state. Which of the following commands will list and remove the correspondent packages?

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

Answer: D

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

NEW QUESTION 161
An administrator attempts to rename a file on a server but receives the following error.

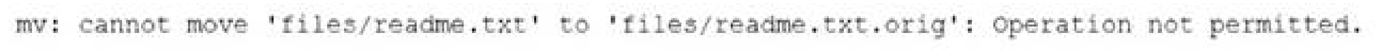

```
mv: cannot move 'files/readme.txt' to 'files/readme.txt.orig': Operation not permitted.
```

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

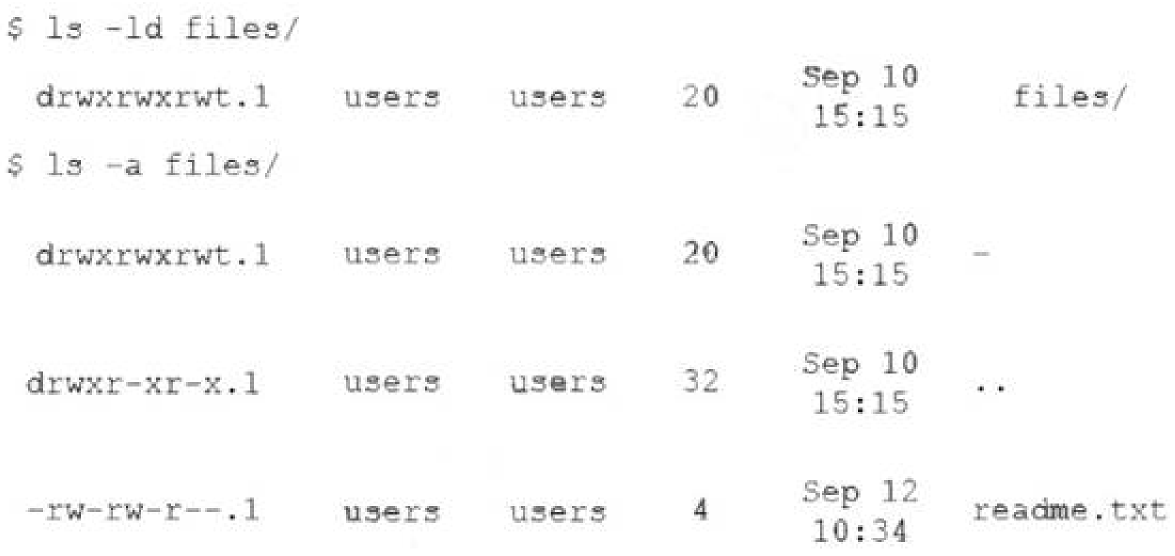

```
$ ls -ld files/
  drwxrwxrwt.1   users   users   20   Sep 10 15:15   files/
$ ls -a files/
  drwxrwxrwt.1   users   users   20   Sep 10 15:15   -
 drwxr-xr-x.1    users   users   32   Sep 10 15:15   ..
 -rw-rw-r--.1    users   users   4    Sep 12 10:34   readme.txt
```

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

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

Answer: D

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

NEW QUESTION 162
A systems administrator created a new Docker image called test. After building the image, the administrator forgot to version the release. Which of the following will allow the administrator to assign the v1 version to the image?

A. docker image save test test:v1
B. docker image build test:vl
C. docker image tag test test:vl
D. docker image version test:v1

Answer: C

Explanation:
The docker image tag test test:v1 command can be used to assign the v1 version to the image called test. This command creates a new tag for the existing image, without changing the original image. The docker image save test test:v1 command would save the image to a file, not assign a version. The docker image build test:vl command is invalid, as vl is not a valid version number. The docker image version test:v1 command does not exist. References: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 16: Virtualization and Cloud Technologies, page 500.

NEW QUESTION 163
A Linux administrator created a new file system. Which of the following files must be updated to ensure the filesystem mounts at boot time?

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

Answer: C

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

NEW QUESTION 164
A cloud engineer needs to launch a container named web-01 in background mode. Which of the following commands will accomplish this task''

A. docker builder -f —name web-01 httpd
B. docker load --name web-01 httpd
C. docker ps -a --name web-01 httpd

D. docker run -d --name web-01 httpd

Answer: D

Explanation:
The docker run -d --name web-01 httpd command will launch a container named web-01 in background mode. This command will create and start a new container from the httpd image, assign it the name web-01, and run it in detached mode (-d), which means the container will run in the background without attaching to the current terminal. The docker builder -f --name web-01 httpd command is invalid, as builder is not a valid docker command, and -f and --name are not valid options for docker build. The docker load --name web-01 httpd command is invalid, as load does not accept a --name option, and httpd is not a valid file name for load. The docker ps -a --name web-01 httpd command is invalid, as ps does not accept a --name option, and httpd is not a valid filter for ps. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Virtualization and Cloud Technologies, page 499.

NEW QUESTION 165
A Linux administrator is reviewing changes to a configuration file that includes the following section:

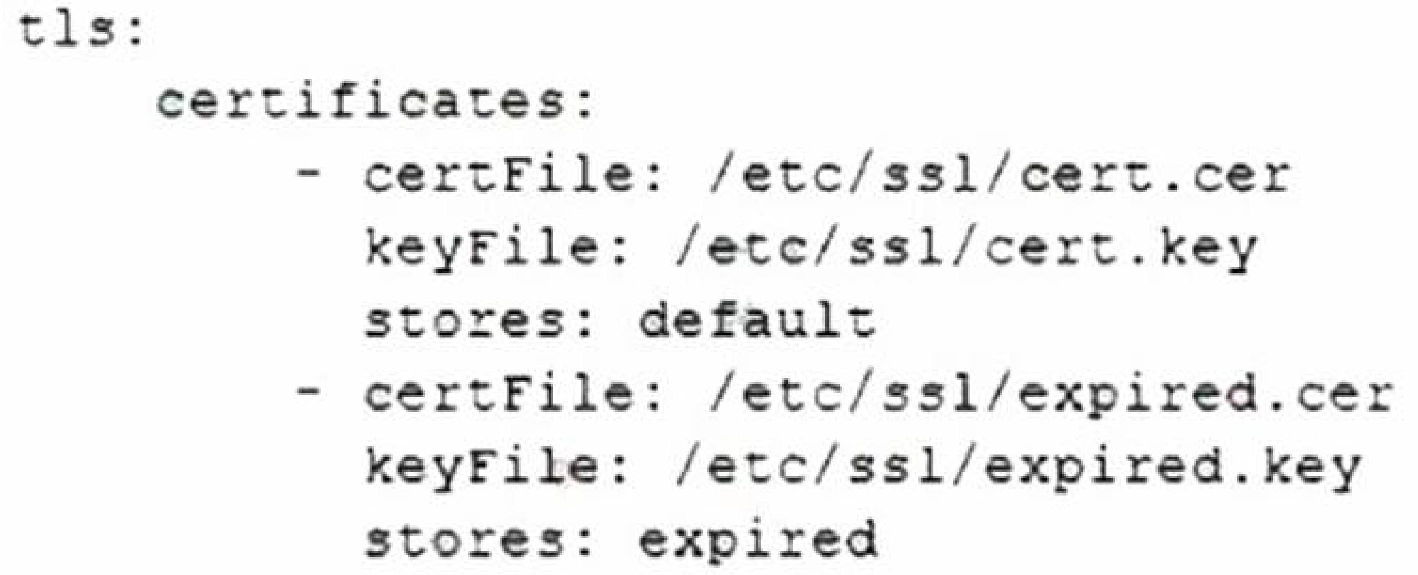

```
tls:
    certificates:
        - certFile: /etc/ssl/cert.cer
          keyFile: /etc/ssl/cert.key
          stores: default
        - certFile: /etc/ssl/expired.cer
          keyFile: /etc/ssl/expired.key
          stores: expired
```

The Linux administrator is trying to select the appropriate syntax formatter to correct any issues with the configuration file. Which of the following should the syntax formatter support to meet this goal?

A. Markdown
B. XML
C. YAML
D. JSON

Answer: C

Explanation:
The configuration file shown in the image is written in YAML format, so the syntax formatter should support YAML to correct any issues with the file. YAML stands for YAML Ain’t Markup Language, and it is a human-readable data serialization language that uses indentation and colons to define key-value pairs. YAML supports various data types, such as scalars, sequences, mappings, anchors, aliases, and tags. The configuration file follows the rules and syntax of YAML, while the other options do not. Markdown is a lightweight markup language that uses plain text formatting to create rich text documents. XML is a markup language that uses tags to enclose elements and attributes. JSON is a data interchange format that uses curly braces to enclose objects and square brackets to enclose arrays. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 21: Automating Tasks with Ansible, page 591.

NEW QUESTION 169
A systems administrator pressed Ctrl+Z after starting a program using the command line, and the shell prompt was presented. In order to go back to the program, which of the following commands can the administrator use?

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

Answer: A

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

NEW QUESTION 171
A systems administrator receives reports that several virtual machines in a host are responding slower than expected. Upon further investigation, the administrator obtains the following output from one of the affected systems:

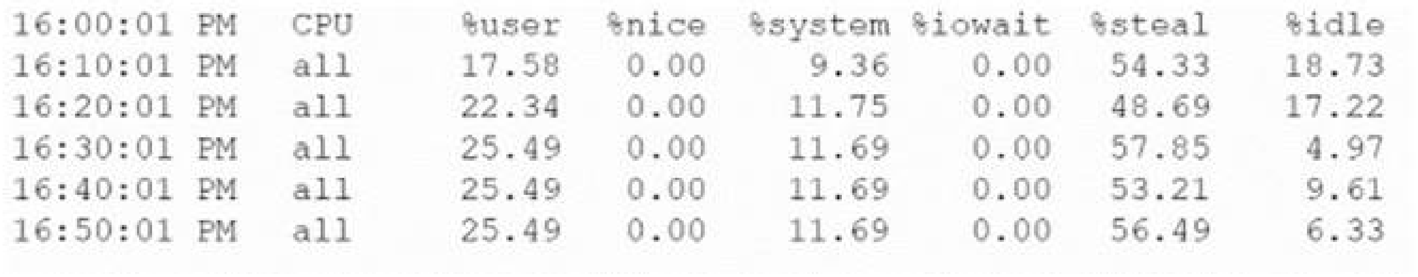

```
16:00:01 PM  CPU     %user  %nice  %system %iowait  %steal   %idle
16:10:01 PM  all     17.58   0.00     9.36    0.00   54.33   18.73
16:20:01 PM  all     22.34   0.00    11.75    0.00   48.69   17.22
16:30:01 PM  all     25.49   0.00    11.69    0.00   57.85    4.97
16:40:01 PM  all     25.49   0.00    11.69    0.00   53.21    9.61
16:50:01 PM  all     25.49   0.00    11.69    0.00   56.49    6.33
```

Which of the following best explains the reported issue?

A. The physical host is running out of CPU resources, leading to insufficient CPU time being allocated to virtual machines.
B. The physical host has enough CPU cores, leading to users running more processes to compensate for the slower response times.
C. The virtual machine has enough CPU cycles, leading to the system use percentage being higher than expected.
D. The virtual machine is running out of CPU resources, leading to users experiencing longer response times.

Answer: D

Explanation:
Based on the output from one of the affected systems, the best explanation for the reported issue is that the virtual machine is running out of CPU resources, leading to users experiencing longer response times (D). The output shows that the system use percentage is very high (57.85%), indicating that the virtual machine is using most of its CPU cycles for system processes. This leaves little CPU time for user processes, which results in slower performance. The other explanations are not supported by the output or are contradictory. References:
? [CompTIA Linux+ Study Guide], Chapter 8: Optimizing Linux Performance, Section: Monitoring CPU Usage
? [How to Interpret CPU Usage Statistics]

NEW QUESTION 173
A Linux administrator has installed a web server, a database server, and a web application on a server. The web application should be active in order to render the web pages. After the administrator restarts the server, the website displays the following message in the browser: Error establishing a database connection. The Linux administrator reviews the following relevant output from the systemd init files:

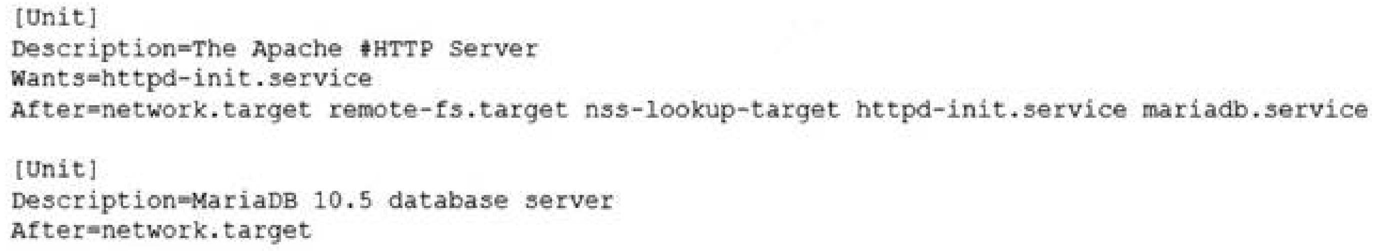

```
[Unit]
Description=The Apache #HTTP Server
Wants=httpd-init.service
After=network.target remote-fs.target nss-lookup-target httpd-init.service mariadb.service

[Unit]
Description=MariaDB 10.5 database server
After=network.target
```

The administrator needs to ensure that the database is available before the web application is started. Which of the following should the administrator add to the HTTP server .service file to accomplish this task?

A. TRIGGERS=mariadb.service
B. ONFAILURE=mariadb.service
C. WANTEDBY=mariadb.service
D. REQUIRES=mariadb.service

Answer: D

Explanation:
The administrator should add REQUIRES=mariadb.service to the HTTP server .service file to ensure that the database is available before the web application is started. This directive specifies that the HTTP server unit requires the MariaDB server unit to be started before it can run. If the MariaDB server unit fails to start or stops for any reason, the HTTP server unit will also fail or stop. This way, the dependency between the web application and the database is enforced by systemd. The other options are not correct directives for accomplishing this task. TRIGGERS=mariadb.service is not a valid directive in systemd unit files. ONFAILURE=mariadb.service means that the HTTP server unit will start only if the MariaDB server unit fails, which is not what we want. WANTEDBY=mariadb.service means that the HTTP server unit will be started when the MariaDB server unit is enabled, but it does not imply a strong dependency or ordering relationship between them. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Services with systemd; systemd.unit(5) - Linux manual page

NEW QUESTION 177
An administrator accidentally installed the httpd RPM package along with several dependencies. Which of the following options is the best way for the administrator to revert the package
installation?

A. dnf clean all
B. rpm -e httpd
C. apt-get clean
D. yum history undo last

Answer: D

Explanation:

The yum history undo last command will undo the last transaction, which in this case is the installation of the httpd RPM package and its dependencies. This will remove the packages that were installed and restore the previous state of the system. See How to undo or redo yum transactions and yum history.References1: https://www.redhat.com/sysadmin/undo- redo-yum-transactions2: https://man7.org/linux/man-pages/man8/yum.8.html#HISTORY

NEW QUESTION 182
......