Exam Questions XK0-005

CompTIA Linux+ Certification Exam

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

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

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

Answer: D

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

NEW QUESTION 3
A Linux user is trying to execute commands with sudo but is receiving the following error:
$ sudo visudo
>>> /etc/sudoers: syntax error near line 28 <<< sudo: parse error in /etc/sudoers near line 28 sudo: no valid sudoers sources found, quitting The following output is provided:
# grep root /etc/shadow root :* LOCK *: 14600 ::::::
Which of the following actions will resolve this issue?

A. Log in directly using the root account and comment out line 28 from /etc/sudoers.
B. Boot the system in single user mode and comment out line 28 from /etc/sudoers.
C. Comment out line 28 from /etc/sudoers and try to use sudo again.
D. Log in to the system using the other regular user, switch to root, and comment out line 28 from /etc/sudoers.

Answer: B

NEW QUESTION 4
A Linux administrator needs to obtain a list of all volumes that are part of a volume group. Which of the following commands should the administrator use to accomplish this task?

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

Answer: B

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

NEW QUESTION 5
A user reported issues when trying to log in to a Linux server. The following outputs were received:

Given the outputs above. which of the following is the reason the user is una-ble to log in to the server?

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

Answer: D

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

NEW QUESTION 6
In which of the following filesystems are system logs commonly stored?

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

Answer: A

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

NEW QUESTION 7
An administrator recently updated the BIND software package and would like to review the default configuration that shipped with this version. Which of the following files should the administrator review?

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

Answer: A

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

NEW QUESTION 8
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 9
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 10
In order to copy data from another VLAN, a systems administrator wants to temporarily assign IP address 10.0.6 5/24 to the newly added network interface enp1s0f1. Which of the following commands should the administrator run to achieve the goal?

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

Answer: A

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

NEW QUESTION 10
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:

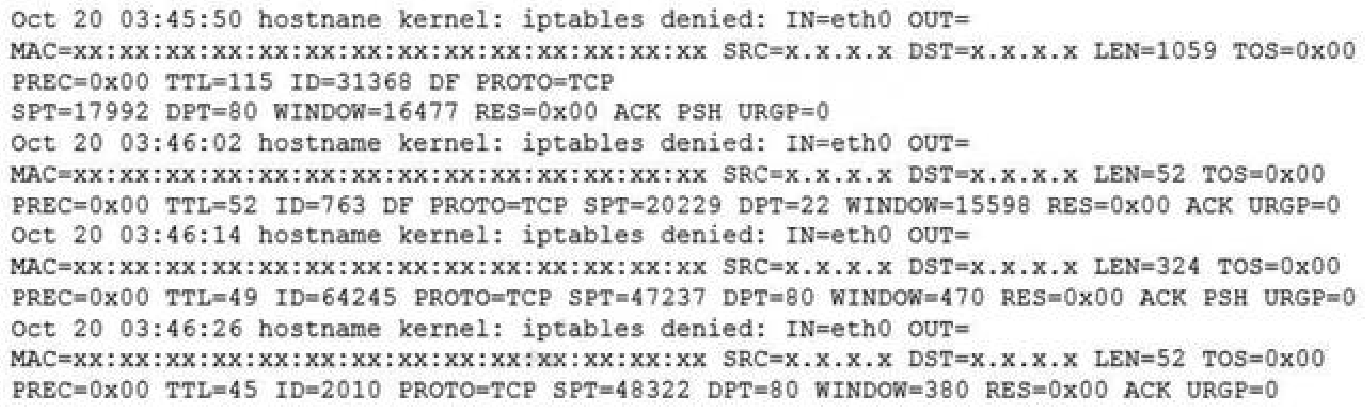

```
Oct 20 03:45:50 hostnane kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx SRC=x.x.x.x DST=x.x.x.x LEN=1059 TOS=0x00
PREC=0x00 TTL=115 ID=31368 DF PROTO=TCP
SPT=17992 DPT=80 WINDOW=16477 RES=0x00 ACK PSH URGP=0
Oct 20 03:46:02 hostname kernel: iptables denied: IN=eth0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

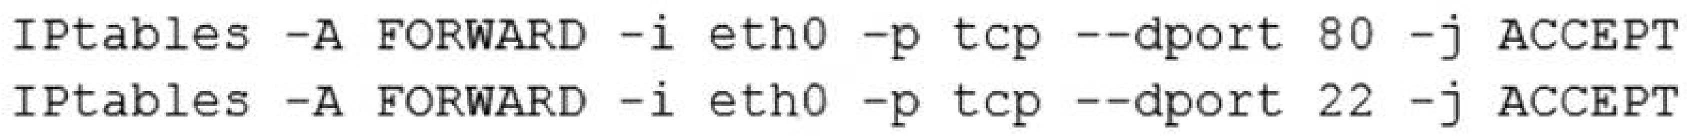

```
IPtables -A FORWARD -i eth0 -p tcp --dport 80 -j ACCEPT
IPtables -A FORWARD -i eth0 -p tcp --dport 22 -j ACCEPT
```

B.

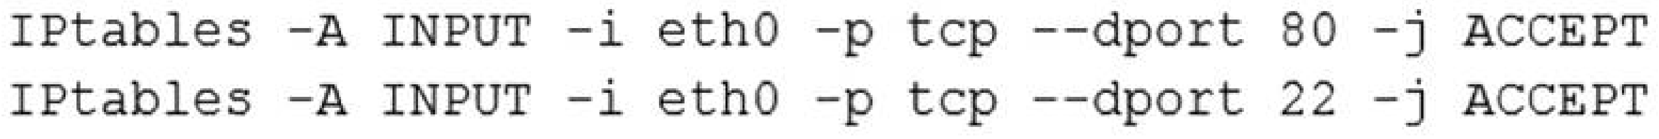

```
IPtables -A INPUT -i eth0 -p tcp --dport 80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --dport 22 -j ACCEPT
```

C.

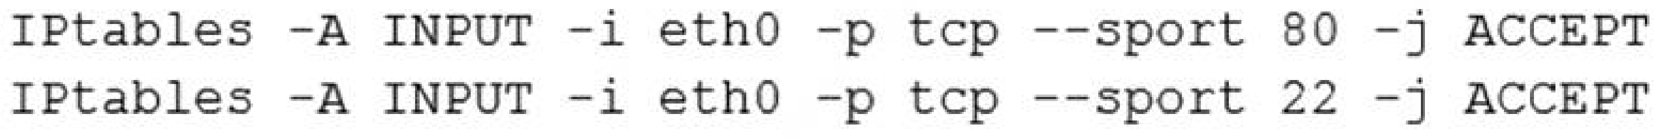

```
IPtables -A INPUT -i eth0 -p tcp --sport 80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --sport 22 -j ACCEPT
```

D.

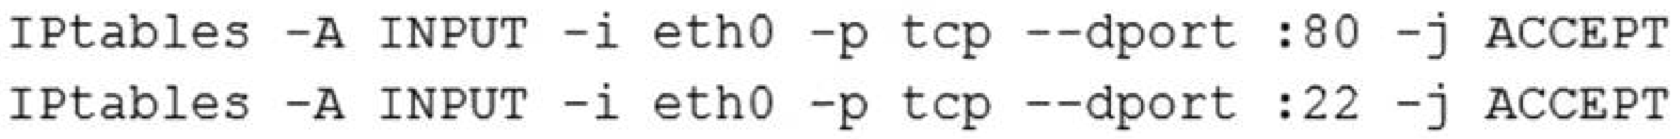

```
IPtables -A INPUT -i eth0 -p tcp --dport :80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --dport :22 -j ACCEPT
```

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 15
A systems administrator needs to clone the partition /dev/sdc1 to /dev/sdd1. Which of the following commands will accomplish this task?

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

Answer: C

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

NEW QUESTION 17
A user generated a pair of private-public keys on a workstation. Which of the following commands will allow the user to upload the public key to a remote server and enable passwordless login?

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

Answer: D

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

NEW QUESTION 18
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 21
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 23
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 25
A Linux administrator is installing a web server and needs to check whether web traffic has already been allowed through the firewall. Which of the following commands should the administrator use to accomplish this task?

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

Answer: C

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

NEW QUESTION 30
A Linux systems administrator is troubleshooting an I/O latency on a single CPU server. The administrator runs a top command and receives the following output:
%Cpu(s): 0.2 us, 33.1 sy, 0.0 ni, 0.0 id, 52.4 wa, 0.0 hi, 0.2 si, 0.0 st
Which of the following is correct based on the output received from the exe-cuted command?

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

Answer: C

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

NEW QUESTION 31
A User on a Linux workstation needs to remotely start an application on a Linux server and then forward the graphical display of that application back to the Linux workstation. Which of the following would enable the user to perform this action?

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

Answer: A

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

NEW QUESTION 32
A cloud engineer needs to change the secure remote login port from 22 to 49000. Which of the following files should the engineer modify to change the port number to the desired value?

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

Answer: D

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

NEW QUESTION 34
A Linux engineer needs to block an incoming connection from the IP address 2.2.2.2 to a secure shell server and ensure the originating IP address receives a response that a firewall is blocking the connection. Which of the following commands can be used to accomplish this task?

A. iptables -A INPUT -p tcp -- dport ssh -s 2.2.2.2 -j DROP
B. iptables -A INPUT -p tcp -- dport ssh -s 2.2.2.2 -j RETURN
C. iptables -A INPUT -p tcp -- dport ssh -s 2.2.2.2 -j REJECT
D. iptables -A INPUT -p tcp -- dport ssh -s 2.2.2.2 -j QUEUE

Answer: C

Explanation:
The REJECT target sends back an error packet to the source IP address, indicating that the connection is refused by the firewall. This is different from the DROP target, which silently discards the packet without any response. The RETURN target returns to the previous chain, which may or may not accept the connection. The QUEUE target passes the packet to a userspace application for further processing, which is not the desired outcome in this case.
References
? CompTIA Linux+ (XK0-005) Certification Study Guide, page 316
? iptables - ssh - access from specific ip only - Server Fault, answer by Eugene Ionichev

NEW QUESTION 37
A Linux system is getting an error indicating the root filesystem is full. Which of the following commands should be used by the systems administrator to resolve this issue? (Choose three.)

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

Answer: ACE

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

NEW QUESTION 41

A systems administrator created a new directory with specific permissions. Given the following output:
# file: comptia
# owner: root
# group: root user: : rwx group :: r-x other: :---
default:user :: rwx default:group :: r-x default:group:wheel: rwx default:mask :: rwx default:other ::-
Which of the following permissions are enforced on /comptia?

A. Members of the wheel group can read files in /comptia.
B. Newly created files in /comptia will have the sticky bit set.
C. Other users can create files in /comptia.
D. Only root can create files in /comptia.

Answer: A

Explanation:

The output shows the file access control list (FACL) of the /comptia directory, which is an extension of the standard Linux permissions that allows more fine-grained control over file and directory access1. The FACL consists of two parts: the access ACL and the default ACL. The access ACL applies to the current object, while the default ACL applies to the objects created within the directory2.
The access ACL has three entries: user, group, and other. These are similar to the standard Linux permissions, but they can be specified for individual users or groups as well. The user entry shows that the owner of the directory (root) has read, write, and execute permissions (rwx). The group entry shows that the group owner of the directory (root) has read and execute permissions (r-x). The other entry shows that all other users have no permissions (—).
The default ACL has five entries: user, group, group:wheel, mask, and other. These are applied to any files or directories created within /comptia. The user entry shows that the owner of the new object will have read, write, and execute permissions (rwx). The group entry shows that the group owner of the new object will have read and execute permissions (r-x). The group:wheel entry shows that the members of the wheel group will have read, write, and execute permissions (rwx) on the new object. The mask entry shows that the maximum permissions allowed for any user or group are read, write, and execute (rwx). The other entry shows that all other users will have no permissions (—) on the new object. Therefore, based on the FACL output, members of the wheel group can read files in /comptia, as they have read permission on both the directory and any files within it. Option B is incorrect because the sticky bit is not set on /comptia or any files within it. The sticky bit is a special permission that prevents users from deleting or renaming files that they do not own in a shared directory3. It is symbolized by a t character in the execute position of others. Option C is incorrect because other users cannot create files in /comptia, as they have no permissions on the directory or any files within it. Option D is incorrect because root is not the only user who can create files in /comptia. Any user who has write permission on the directory can create files within it, such as members of the wheel group.

NEW QUESTION 42

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

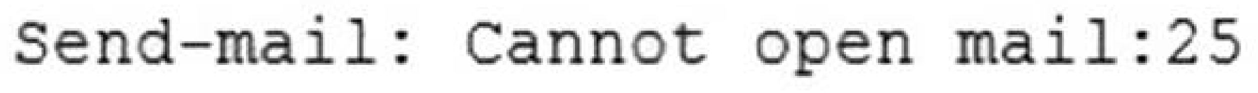

```
Send-mail: Cannot open mail:25
```

The local machine DNS settings are:

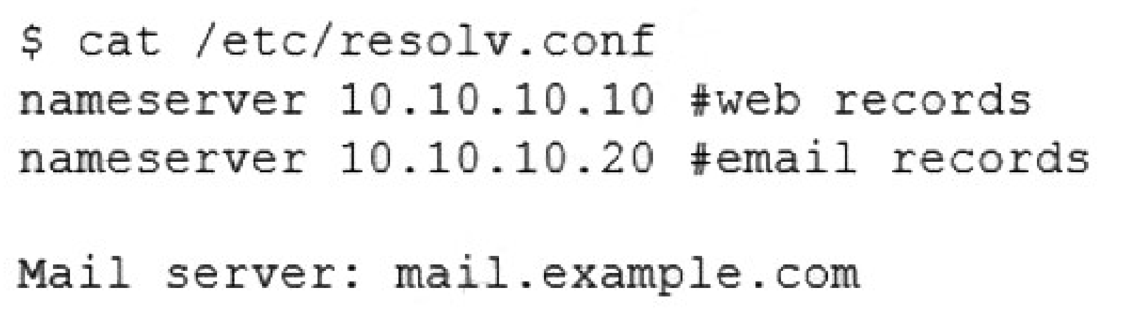

```
$ cat /etc/resolv.conf
nameserver 10.10.10.10 #web records
nameserver 10.10.10.20 #email records

Mail server: mail.example.com
```

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

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

Answer: B

Explanation:

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

NEW QUESTION 47

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

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

Answer: B

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

NEW QUESTION 48
A Linux administrator is trying to remove the ACL from the file /home/user/data. txt but receives the following error message:

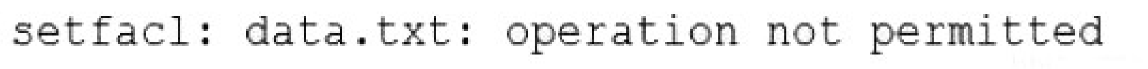

```
setfacl: data.txt: operation not permitted
```

Given the following analysis:

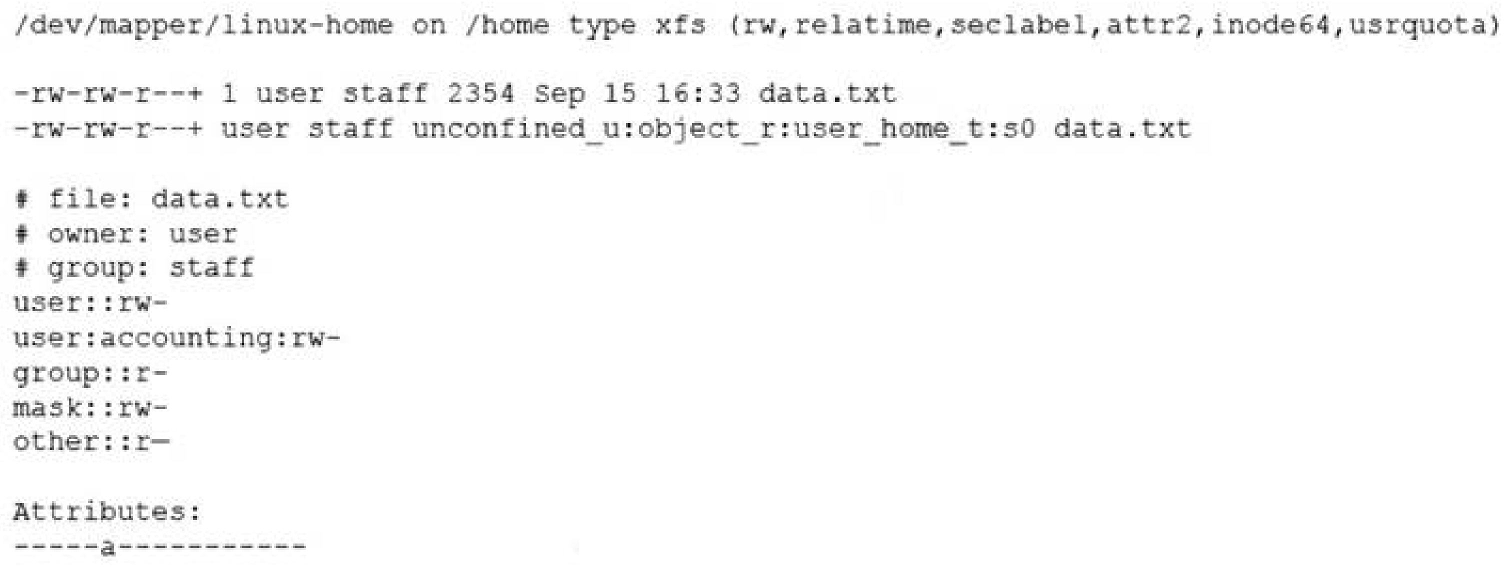

```
/dev/mapper/linux-home on /home type xfs (rw,relatime,seclabel,attr2,inode64,usrquota)

-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 50
Users in the human resources department are trying to access files in a newly created directory. Which of the following commands will allow the users access to the files?

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

Answer: B

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

NEW QUESTION 55
A systems administrator is enabling LUKS on a USB storage device with an ext4 filesystem format. The administrator runs dmesg and notices the following output:

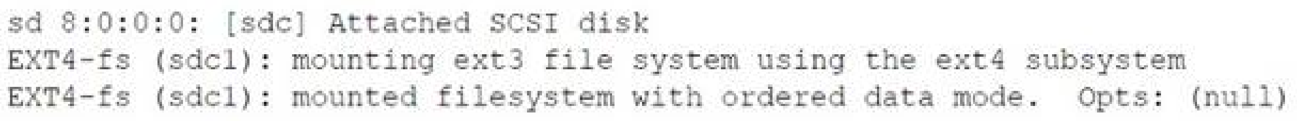

```
sd 8:0:0:0: [sdc] Attached SCSI disk
EXT4-fs (sdc1): mounting ext3 file system using the ext4 subsystem
EXT4-fs (sdc1): mounted filesystem with ordered data mode.  Opts: (null)
```

Given this scenario, which of the following should the administrator perform to meet these requirements? (Select three).

A. gpg /dev/sdcl
B. pvcreate /dev/sdc
C. mkfs . ext4 /dev/mapper/LUKSCJ001 - L ENCRYPTED
D. umount / dev/ sdc
E. fdisk /dev/sdc
F. mkfs . vfat /dev/mapper/LUKS0001 — L ENCRYPTED
G. wipefs —a/dev/sdbl
H. cryptsetup IuksFormat /dev/ sdcl

Answer: CDH

Explanation:
To enable LUKS on a USB storage device with an ext4 filesystem format, the administrator needs to perform the following steps:
? Unmount the device if it is mounted using umount /dev/sdc (D)
? Create a partition table on the device using fdisk /dev/sdc (E)
? Format the partition with LUKS encryption using cryptsetup luksFormat /dev/sdc1 (H)
? Open the encrypted partition using cryptsetup luksOpen /dev/sdc1 LUKS0001
? Create an ext4 filesystem on the encrypted partition using mkfs.ext4 /dev/mapper/LUKS0001 ©
? Mount the encrypted partition using mount /dev/mapper/LUKS0001 /mnt References:
? [CompTIA Linux+ Study Guide], Chapter 9: Securing Linux, Section: Encrypting Disks
? [How to Encrypt USB Drive on Ubuntu 18.04]

NEW QUESTION 57
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.

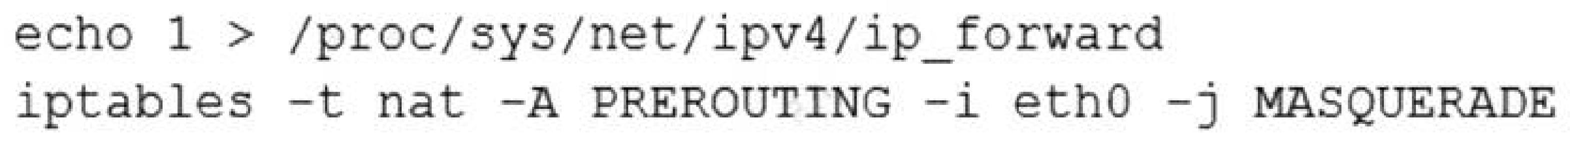

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A PREROUTING -i eth0 -j MASQUERADE
```

A.

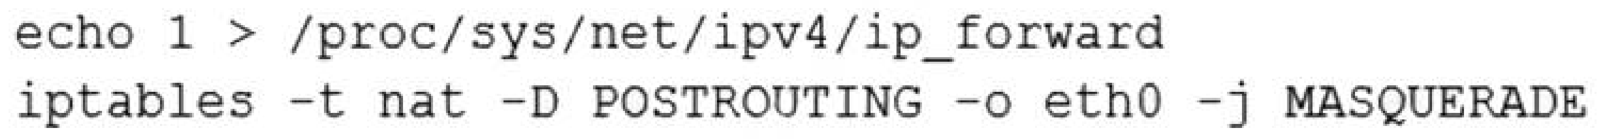

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -D POSTROUTING -o eth0 -j MASQUERADE
```

B.

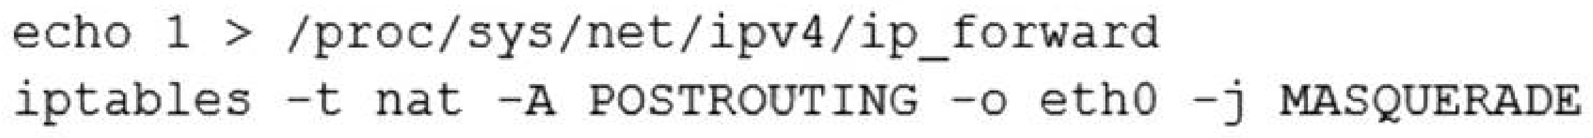

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
```

C.

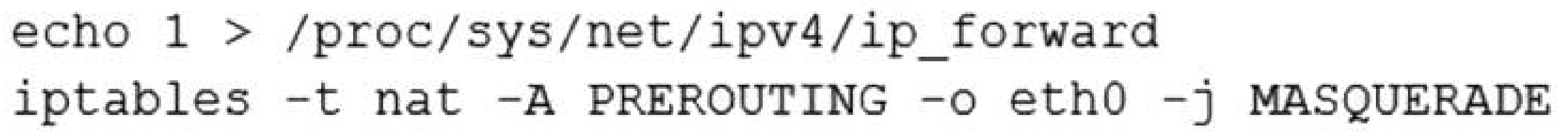

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A PREROUTING -o eth0 -j MASQUERADE
```

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 60
Following the migration from a disaster recovery site, a systems administrator wants a server to require a user to change credentials at initial login. Which of the following commands should
be used to ensure the aging attribute?

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

Answer: B

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

NEW QUESTION 64
A Linux administrator found many containers in an exited state. Which of the following commands will allow the administrator to clean up the containers in an exited state?

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

Answer: B

Explanation:
This command will remove all containers, regardless of their state, by passing the IDs of all containers to the docker rm command. The docker ps -aq command will list the IDs of all containers, including the ones in an exited state, and the $ ( ) syntax will substitute the output of the command as an argument for the docker rm command. This is a quick and easy way to clean up all containers, but it may also remove containers that are still needed or running.
References
? docker rm | Docker Docs - Docker Documentation, section “Remove all containers”
? Docker Remove Exited Containers | Easy methods. - Bobcares, section “For removing all exited containers”

NEW QUESTION 69
An administrator is trying to diagnose a performance issue and is reviewing the following output:

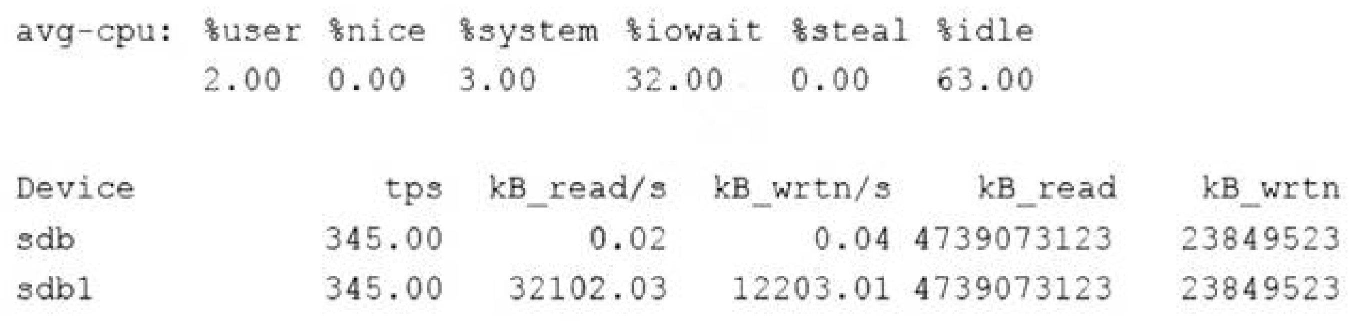

```
avg-cpu: %user  %nice  %system  %iowait  %steal  %idle
          2.00   0.00   3.00     32.00    0.00    63.00

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 70
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 74
A Linux administrator is adding a new configuration file to a Git repository. Which of the following describes the correct order of Git commands to accomplish the task successfully?

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

Answer: B

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

NEW QUESTION 78
Due to performance issues on a server, a Linux administrator needs to termi-nate an unresponsive process. Which of the following commands should the administrator use to terminate the process immediately without waiting for a graceful shutdown?

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

Answer: A

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

NEW QUESTION 82
A systems administrator is adding a Linux-based server and removing a Windows-based server from a cloud-based environment. The changes need to be validated before they are applied to the cloud-based environment. Which of the following tools should be used to meet this requirement?

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

Answer: D

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

NEW QUESTION 84
A Linux administrator needs to determine whether a hostname is in the DNS. Which of the following would supply the information that is needed?

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

Answer: A

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

NEW QUESTION 88
Users are unable to create new files on the company's FTP server, and an administrator is troubleshooting the issue. The administrator runs the following commands:

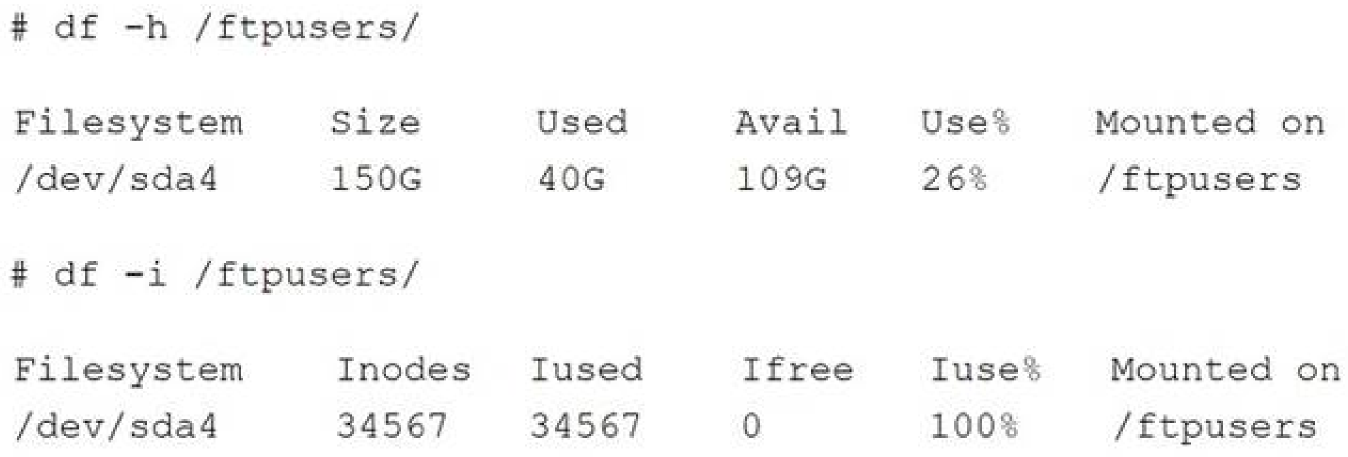

```
# df -h /ftpusers/

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

# df -i /ftpusers/

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

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

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

Answer: C

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

NEW QUESTION 92
A Linux administrator has set up a new DNS forwarder and is configuring all internal servers to use the new forwarder to look up external DNS requests. The administrator needs to modify the firewall on the server for the DNS forwarder to allow the internal servers to communicate to it and make the changes persistent between server reboots. Which of the following commands should be run on the DNS forwarder server to accomplish this task?

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

Answer: D

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

NEW QUESTION 94
A systems administrator was tasked with assigning the temporary IP address/netmask 192.168.168.1/255.255.255.255 to the interface eth0 of a Linux server.
When adding the address, the following error appears:
# ip address add 192.168.168.1/33 dev eth0
Error: any valid prefix is expected rather than "192.168.168.1/33".
Based on the command and its output above, which of the following is the cause of the issue?

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

Answer: A

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

NEW QUESTION 99
A Linux system is failing to boot with the following error:

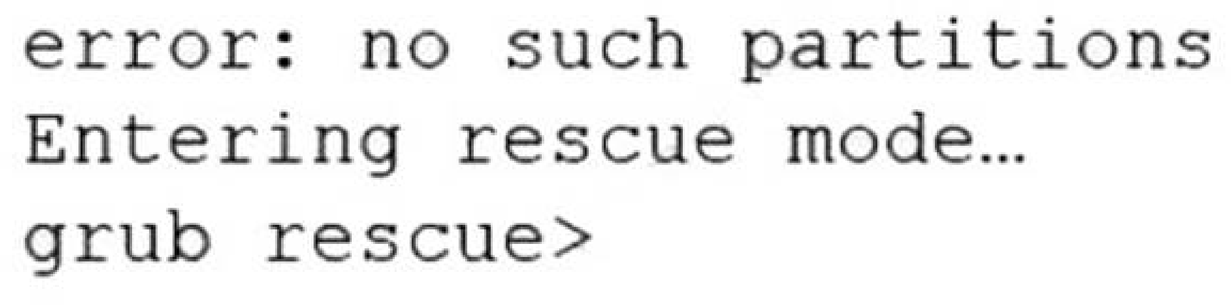

```
error: no such partitions
Entering rescue mode...
grub rescue>
```

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

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

Answer: BF

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

NEW QUESTION 103
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 106
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 107
A systems administrator received a request to change a user's credentials. Which of the following commands will grant the request?

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

Answer: A

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

NEW QUESTION 109
A Linux administrator found many containers in an exited state. Which of the following commands will allow the administrator to clean up the containers in an exited state?

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

Answer: B

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

NEW QUESTION 110
An administrator would like to list all current containers, regardless of their running state. Which of the following commands would allow the administrator to accomplish this task?

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

Answer: A

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

NEW QUESTION 113
A Linux system is having issues. Given the following outputs:
# dig @192.168.2.2 mycomptiahost
; << >> DiG 9.9.4-RedHat-9.9.4-74.el7_6.1 << >> @192.168.2.2 mycomptiahost
; (1 server found)
;; global options: +cmd
;; connection timed out; no servers could be reached
# nc -v 192.168.2.2 53
Ncat: Version 7.70 (https://nmap.org/ncat) Ncat: Connection timed out.
# ping 192.168.2.2
PING 192.168.2.2 (192.168.2.2) 56(84) bytes of data.
64 bytes from 192.168.2.2: icmp_seq=1 ttl=117 time=4.94 ms 64 bytes from 192.168.2.2: icmp_seq=2 ttl=117 time=10.5 ms Which of the following best describes this issue?

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

Answer: D

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

NEW QUESTION 114
A Linux administrator is tasked with adding users to the system. However, the administrator wants to ensure the users' access will be disabled once the project is over. The expiration date should be 2021-09-30. Which of the following commands will accomplish this task?

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

Answer: A

Explanation:
The command that will accomplish this task is sudo useradd -e 2021-09-30 Project_user. This command will create a new user account named Project_user with an expiration date of 2021-09-30. The -e option of useradd specifies the date on which the user account will be disabled in YYYY-MM-DD format.
The other options are not correct commands for creating a user account with an expiration date. The sudo useradd -c 2021-09-30 Project_user command will create a new user account named Project_user with a comment of 2021-09-30. The -c option of useradd specifies a comment or description for the user account, not an expiration date. The sudo modinfo -F 2021-09-30 Project_user command is invalid because modinfo is not a command for managing user accounts, but a command for displaying information about kernel modules. The -F option of modinfo specifies a field name to show, not an expiration date. The sudo useradd -m -d 2021-09-30 Project_user command will create a new user account named Project_user with a home directory of 2021-09-30. The -m option of useradd specifies that the home directory should be created if it does not exist, and the -d option specifies the home directory name, not an expiration date. References: useradd(8) - Linux manual page; modinfo(8) - Linux manual page

NEW QUESTION 116
Employees in the finance department are having trouble accessing the file /opt/work/file. All IT employees can read and write the file. Systems administrator reviews the following output:

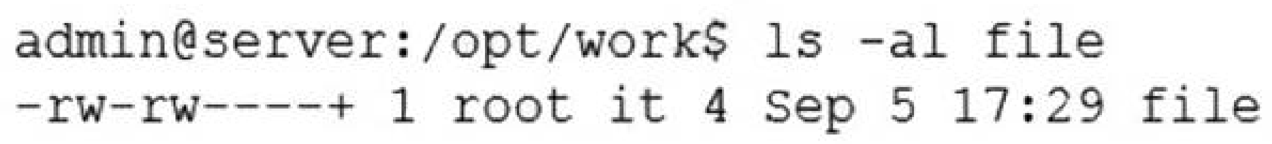

```
admin@server:/opt/work$ ls -al file
-rw-rw----+ 1 root it 4 Sep 5 17:29 file
```

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

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

Answer: D

Explanation:
The command setfacl -m g:finance:rw file will permanently fix the access issue while limiting access to IT and finance department employees. The setfacl command is a tool for modifying the access control lists (ACLs) of files and directories on Linux systems. The ACLs are a mechanism that allows more fine-grained control over the permissions of files and directories than the traditional owner-group-others model. The - m option specifies the modification to the ACL. The

g:finance:rw means that the group named finance will have read and write permissions on the file. The file is the name of the file to modify, in this case /opt/work/file. The command setfacl -m g:finance:rw file will add an entry to the ACL of the file that will grant read and write access to the finance group.
This will fix the access issue and allow the finance employees to access the file. The command will also preserve the existing permissions of the file, which means that the IT employees will still have read and write access to the file. This will limit the access to IT and finance department employees and prevent unauthorized access from other users.
This is the correct command to use to accomplish the task. The other options are incorrect because they either do not fix the access issue (chattr +i file or chown it:finance file) or do not limit the access to IT and finance department employees (chmod 666 file). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing File Permissions and Ownership, page 352.

NEW QUESTION 120
A systems administrator is tasked with creating a cloud-based server with a public IP address.

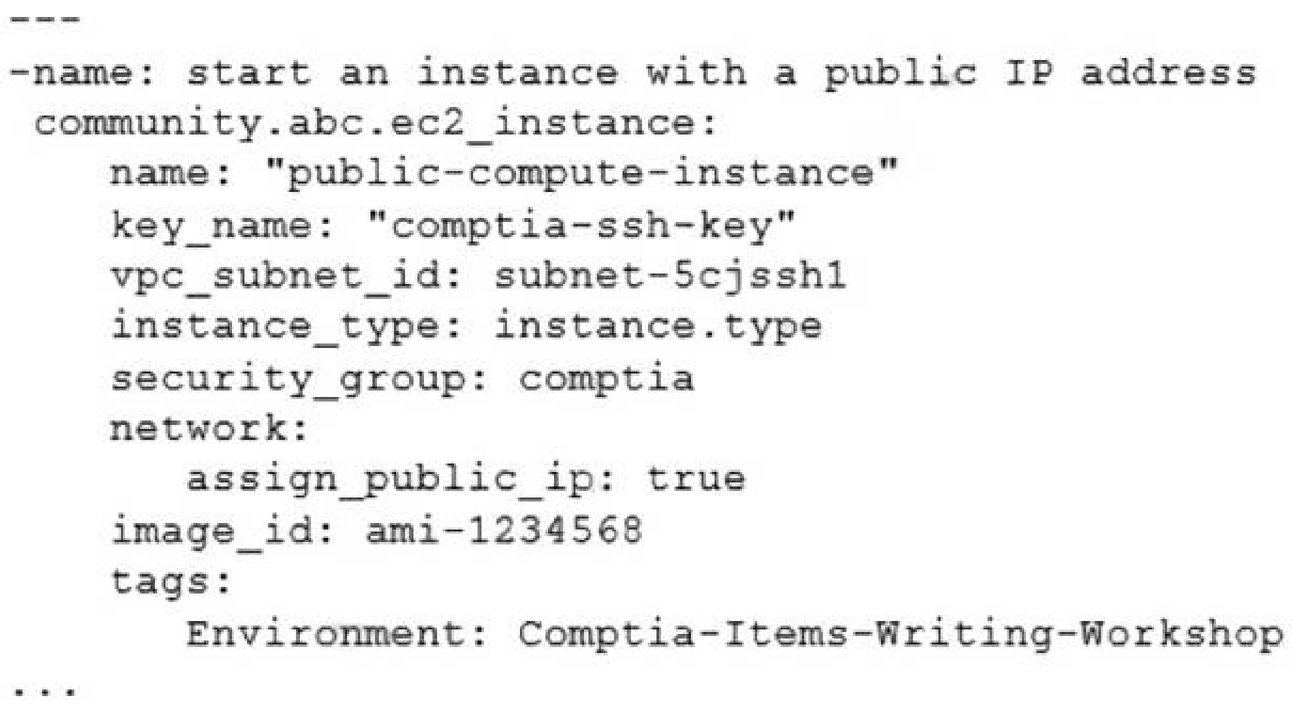

```
---
-name: start an instance with a public IP address
 community.abc.ec2_instance:
    name: "public-compute-instance"
    key_name: "comptia-ssh-key"
    vpc_subnet_id: subnet-5cjssh1
    instance_type: instance.type
    security_group: comptia
    network:
       assign_public_ip: true
    image_id: 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 122
A systems administrator is implementing a new service task with systems at startup and needs to execute a script entitled test.sh with the following content:

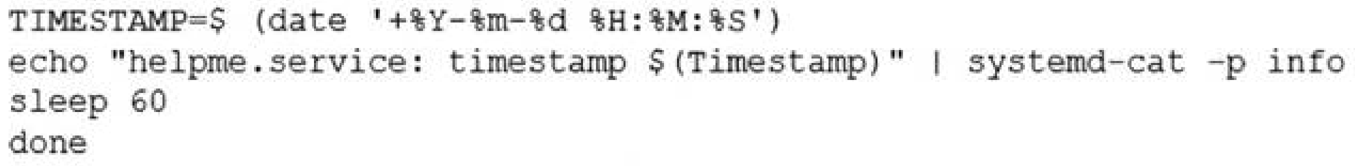

```
TIMESTAMP=$ (date '+%Y-%m-%d %H:%M:%S')
echo "helpme.service: timestamp $(Timestamp)" | systemd-cat -p info
sleep 60
done
```

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

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

Answer: BC

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

NEW QUESTION 127
A Linux administrator has defined a systemd script docker-repository.mount to mount a volume for use by the Docker service. The administrator wants to ensure that Docker service does not start until the volume is mounted. Which of the following configurations needs to be added to the Docker service definition to best accomplish this task?

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

Answer: C

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

NEW QUESTION 128
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 129
Which of the following specifications is used to perform disk encryption in a Linux system?

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

Answer: A

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

NEW QUESTION 130
A Linux administrator is trying to start the database service on a Linux server but is not able to run it. The administrator executes a few commands and receives the following output:

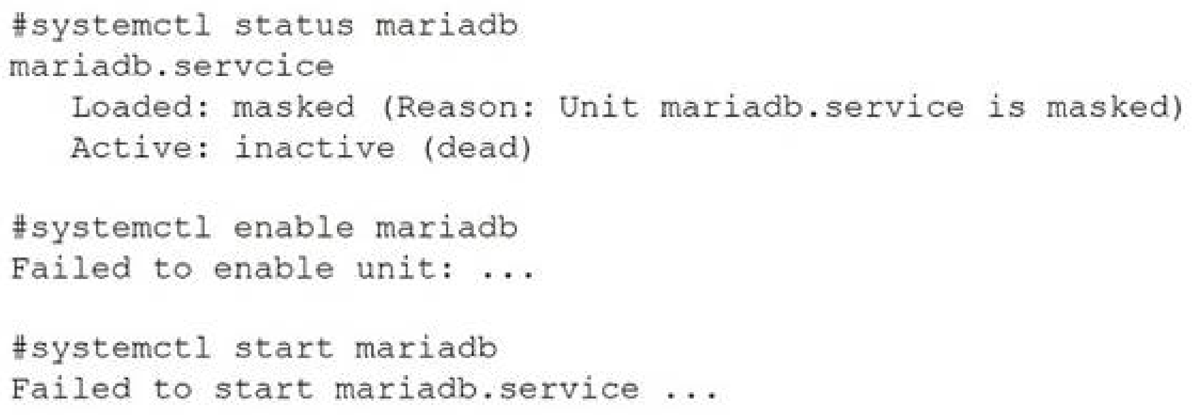

```
#systemctl status mariadb
mariadb.servcice
   Loaded: masked (Reason: Unit mariadb.service is masked)
   Active: inactive (dead)

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

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

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

A. systemctl unmask mariadb
B. journalctl —g mariadb
C. dnf reinstall mariadb
D. systemctl start mariadb

E. chkconfig mariadb on
F. service mariadb reload

Answer: AD

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

NEW QUESTION 133
A database administrator requested the installation of a custom database on one of the servers. Which of the following should the Linux administrator configure so the requested packages can be installed?

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

Answer: C

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

NEW QUESTION 137
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:

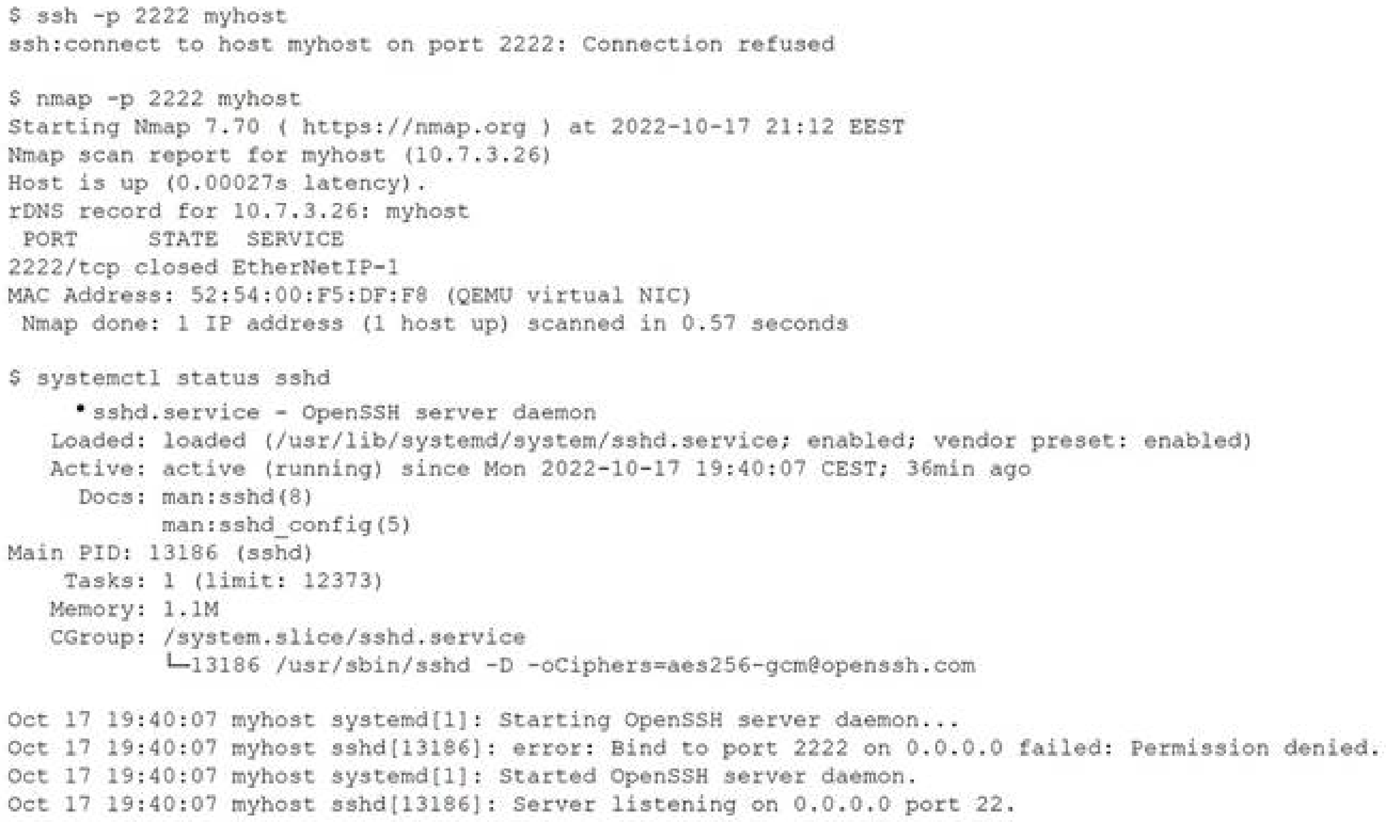

```
$ ssh -p 2222 myhost
ssh:connect to host myhost on port 2222: Connection refused

$ 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 141
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 145
At what point is the Internal Certificate Authority (ICA) created?

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

Answer: A

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

NEW QUESTION 147
A Linux administrator rebooted a server. Users then reported some of their files were missing. After doing some troubleshooting, the administrator found one of the filesystems was missing. The filesystem was not listed in /etc/f stab and might have been mounted manually by someone prior to reboot. Which of the following would prevent this issue from reoccurring in the future?

A. Sync the mount units.
B. Mount the filesystem manually.
C. Create a mount unit and enable it to be started at boot.
D. Remount all the missing filesystems

Answer: C

Explanation:
The best way to prevent this issue from reoccurring in the future is to create a mount unit and enable it to be started at boot. A mount unit is a systemd unit that defines how and where a filesystem should be mounted. By creating a mount unit for the missing filesystem and enabling it with systemct1 enable, the administrator can ensure that the filesystem will be automatically mounted at boot time, regardless of whether it is listed in /etc/fstab or not. Syncing the mount units will not prevent the issue, as it will only synchronize the state of existing mount units with /etc/fstab, not create new ones. Mounting the filesystem manually will not prevent the issue, as it will only mount the filesystem temporarily, not permanently. Remounting all the missing filesystems will not prevent the issue, as it will only mount the filesystems until the next reboot, not after. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Disk Storage, page 457.

NEW QUESTION 149
A systems administrator is tasked with creating an Ansible playbook to automate the installation of patches on several Linux systems. In which of the following languages should the playbook be written?

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

Answer: B

Explanation:
The language that the playbook should be written in is YAML. YAML stands for YAML Ain't Markup Language, which is a human-readable data serialization language. YAML is commonly used for configuration files and data exchange. YAML uses indentation, colons, dashes, and brackets to represent the structure and values of the data. YAML also supports comments, variables, expressions, and functions. Ansible is an open-source tool for automating tasks and managing configuration on Linux systems. Ansible uses YAML to write playbooks, which are files that define the desired state and actions for the systems. Playbooks can be used to automate the installation of patches on several Linux systems by specifying the hosts, tasks, modules, and parameters. The language that the playbook

should be written in is YAML. This is the correct answer to the question. The other options are incorrect because they are not the languages that Ansible uses for playbooks (SQL, HTML, or JSON). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 549.

NEW QUESTION 152
A systems administrator has been tasked with disabling the nginx service from the environment to prevent it from being automatically and manually started. Which of the following commands will accomplish this task?

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

Answer: C

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

NEW QUESTION 155
A systems administrator configured firewall rules using firewalld. However, after the system is rebooted, the firewall rules are not present:

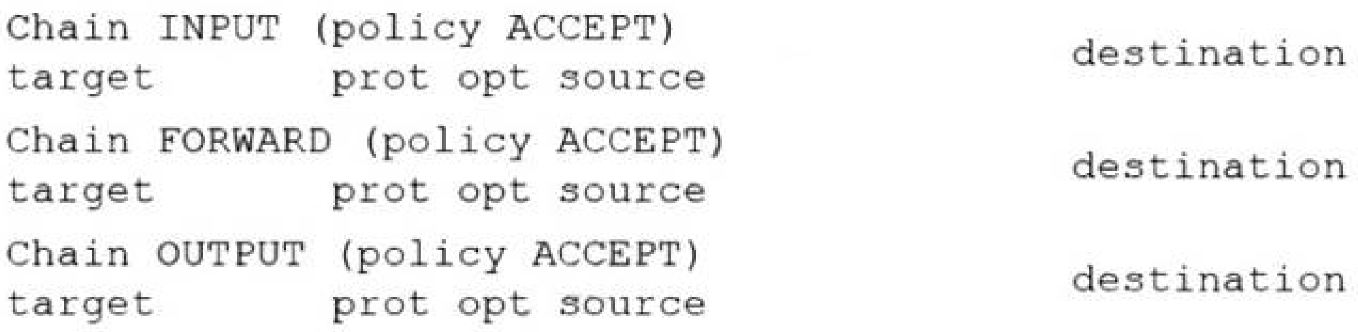

```
Chain INPUT (policy ACCEPT)
target         prot opt source                 destination
Chain FORWARD (policy ACCEPT)
target         prot opt source                 destination
Chain OUTPUT (policy ACCEPT)
target         prot opt source                 destination
```

The systems administrator makes additional checks:

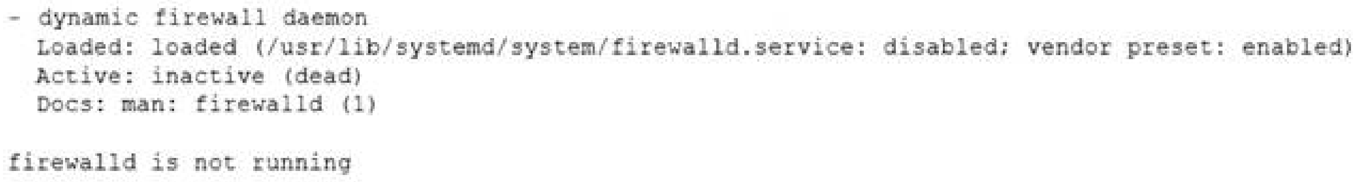

```
- dynamic firewall daemon
  Loaded: loaded (/usr/lib/systemd/system/firewalld.service: disabled; vendor preset: enabled)
  Active: inactive (dead)
  Docs: man: firewalld (1)

firewalld is not running
```

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

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

Answer: D

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

NEW QUESTION 157
The security team has identified a web service that is running with elevated privileges A Linux administrator is working to change the systemd service file to meet security compliance standards. Given the following output:

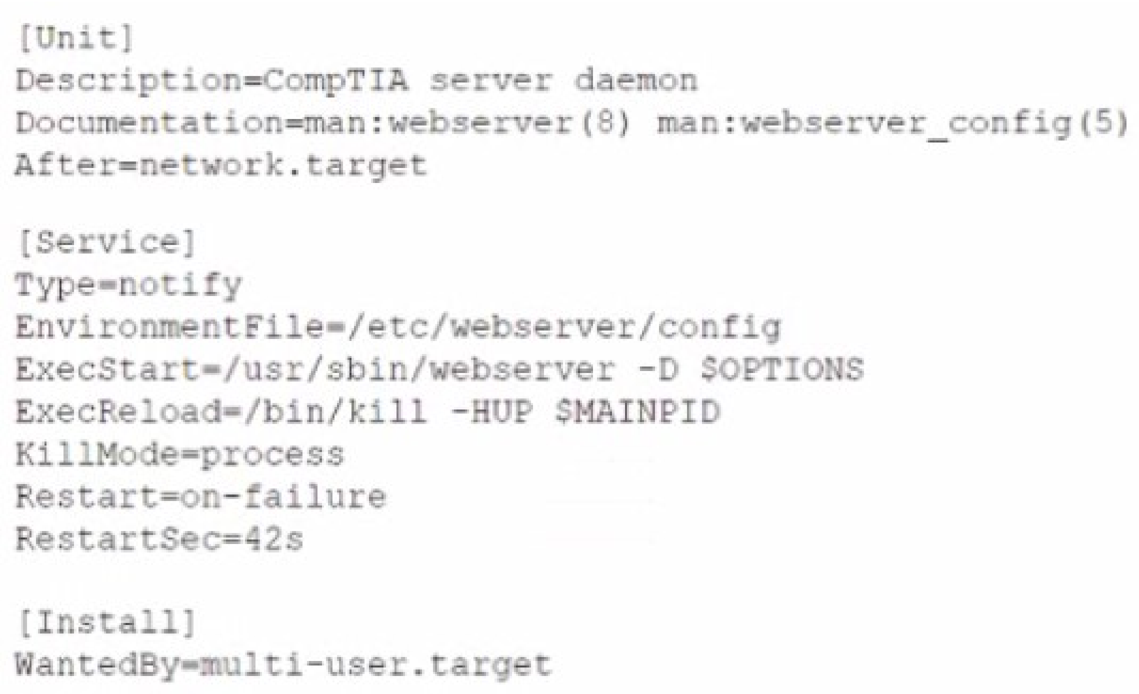

```
[Unit]
Description=CompTIA server daemon
Documentation=man:webserver(8) man:webserver_config(5)
After=network.target

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

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

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

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

Answer: C

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

NEW QUESTION 160
A Linux administrator needs to resolve a service that has failed to start. The administrator runs the following command:

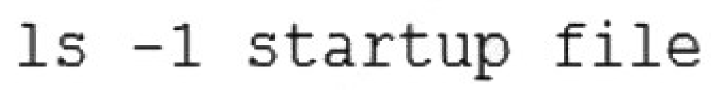

```
ls -1 startup file
```

The following output is returned

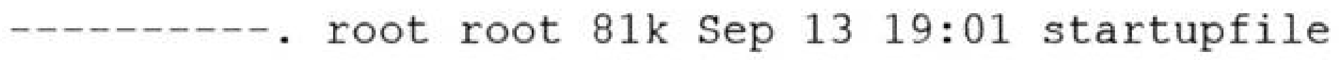

```
----------. root root 81k Sep 13 19:01 startupfile
```

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 163
A systems administrator wants to list all local accounts in which the UID is greater than 500. Which of the following commands will give the correct output?

A. find /etc/passwd —size +500
B. cut —d: fl / etc/ passwd > 500
C. awk -F: ‘$3 > 500 {print $1}' /etc/passwd
D. sed '/UID/' /etc/passwd < 500

Answer: C

Explanation:
The correct command to list all local accounts in which the UID is greater than 500 is:

awk -F: '$3 > 500 {print $1}' /etc/passwd
This command uses awk to process the /etc/passwd file, which contains information about the local users on the system. The -F: option specifies that the fields are separated by colons. The $3 refers to the third field, which is the UID. The condition $3 > 500 filters out the users whose UID is greater than 500. The action {print $1} prints the first field, which is the username.
The other commands are incorrect because:
? find /etc/passwd —size +500 will search for files that are larger than 500 blocks in size, not users with UID greater than 500.
? cut —d: fl / etc/ passwd > 500 will cut the first field of the /etc/passwd file using colon as the delimiter, but it will not filter by UID or print only the usernames. The > 500 part will redirect the output to a file named 500, not compare with the UID.
? sed '/UID/' /etc/passwd < 500 will use sed to edit the /etc/passwd file and replace any line that contains UID with 500, not list the users with UID greater than 500. The < 500 part will redirect the input from a file named 500, not compare with the UID.
References:
? Linux List All Users In The System Command - nixCraft, section “List all users in Linux using /etc/passwd file”.
? Unix script getting users with UID bigger than 500 - Stack Overflow, section “Using awk”.

NEW QUESTION 165
A systems administrator intends to use a UI-JID to mount a new partition per-manently on a Linux system. Which of the following commands can the adminis-trator run to obtain information about the UUlDs of all disks attached to a Linux system?

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

Answer: B

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

NEW QUESTION 168
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 172
A Linux administrator needs to create a symlink for /usr/local/bin/app-a, which was installed in /usr/local/share/app-a. Which of the following commands should the administrator use?

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

Answer: A

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

NEW QUESTION 176
A systems administrator is troubleshooting connectivity issues and trying to find out why a Linux server is not able to reach other servers on the same subnet it is connected to. When listing link parameters, the following is presented:

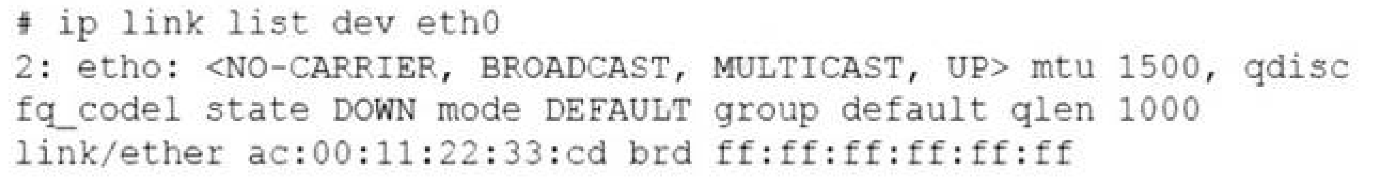

```
# ip link list dev eth0
2: etho: <NO-CARRIER, BROADCAST, MULTICAST, UP> mtu 1500, qdisc
fq_codel state DOWN mode DEFAULT group default qlen 1000
link/ether ac:00:11:22:33:cd brd ff:ff:ff:ff:ff:ff
```

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

A. The address ac:00:11:22:33:cd is not a valid Ethernet address.
B. The Ethernet broadcast address should be ac:00:11:22:33:ff instead.
C. The network interface eth0 is using an old kernel module.
D. The network interface cable is not connected to a switch.

Answer: D

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

NEW QUESTION 181
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:

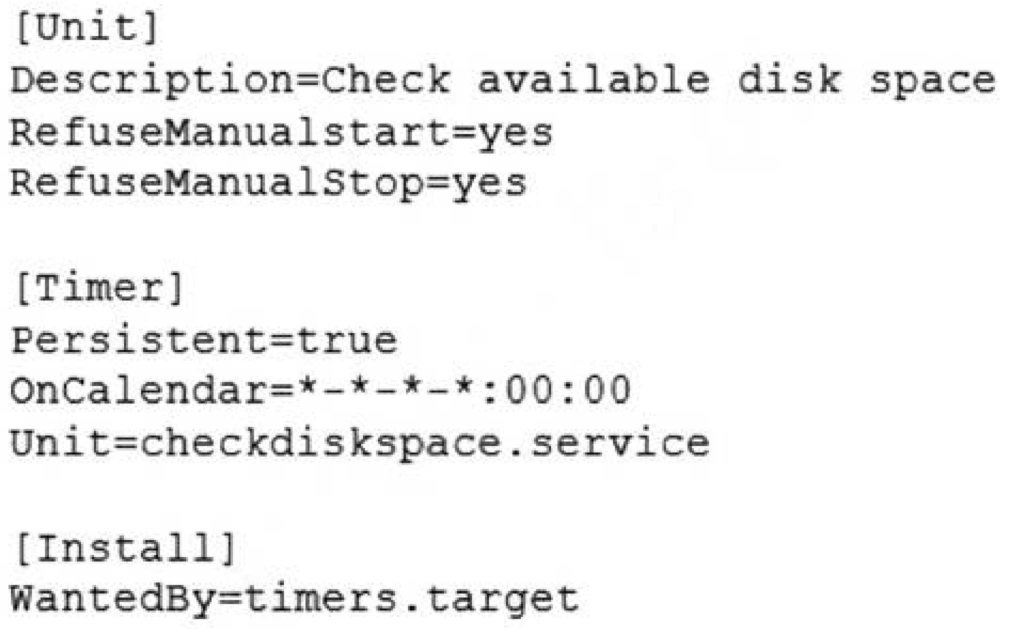

```
[Unit]
Description=Check available disk space
RefuseManualstart=yes
RefuseManualStop=yes

[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:

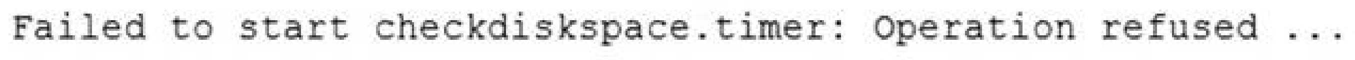

```
Failed to start checkdiskspace.timer: Operation refused ...
```

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 184
A cloud engineer wants to delete all unused networks that are not referenced by any container. Which of the following commands will achieve this goal?

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

Answer: C

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

NEW QUESTION 186
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 188
Joe, a user, is unable to log in to the Linux system. Given the following output:

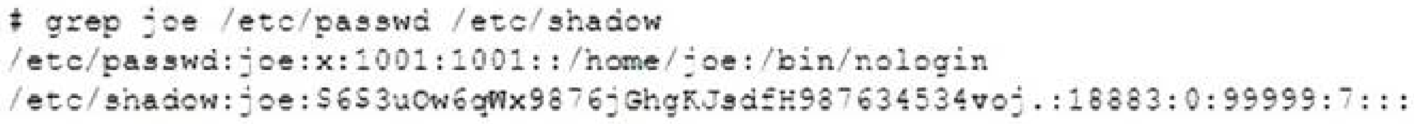

```
# grep joe /etc/passwd /etc/shadow
/etc/passwd:joe:x:1001:1001::/home/joe:/bin/nologin
/etc/shadow:joe:$6$3uOw6qWx9876jGhgKJsdfH987634534voj.:18883:0:99999:7:::
```

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 193
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 197
A Linux engineer finds multiple failed login entries in the security log file for application users. The Linux engineer performs a security audit and discovers a security issue. Given the following:
# grep -iE '*www*|db' /etc/passwd
www-data:x:502:502:www-data:/var/www:/bin/bash db:x: 505:505:db: /opt/db:/bin/bash
Which of the following commands would resolve the security issue?

A. usermod -d /srv/www-data www-data && usermod -d /var/lib/db db
B. passwd -u www-data && passwd -u db
C. renice -n 1002 -u 502 && renice -n 1005 -u 505
D. chsh -s /bin/false www-data && chsh -s /bin/false db

Answer: D

Explanation:
This command will use the chsh tool to change the login shell of the users www-data and db to /bin/false, which means they will not be able to log in to the system1. This will prevent unauthorized access attempts and improve security.
References: 1: Replacing /bin/bash with /bin/false in /etc/passwd file

NEW QUESTION 198
A developer is trying to install an application remotely that requires a graphical interface for installation. The developer requested assistance to set up the necessary environment variables along with X11 forwarding in SSH. Which of the following environment variables must be set in remote shell in order to launch the graphical interface?

A. $RHOST
B. SETENV
C. $SHELL
D. $DISPLAY

Answer: D

Explanation:
The environment variable that must be set in remote shell in order to launch the graphical interface is $DISPLAY. This variable tells X11 applications where to display their windows on screen. It usually has the form hostname:displaynumber.screennumber, where hostname is the name of the computer running the X server, displaynumber is a unique identifier for an X display on that computer, and screennumber is an optional identifier for a screen within an X display. For example, localhost:0.0 means display number 0 on the local host. If the hostname is omitted, it defaults to the local host.
The other options are not correct environment variables for launching the graphical interface. $RHOST is a variable that stores the name of the remote host, but it is not used by X11 applications. SETENV is a command that sets environment variables in some shells, but it is not an environment variable itself. $SHELL is a variable that stores the name of the current shell, but it is not related to X11 forwarding. References: How to enable or disable X11 forwarding in an SSH server; How to Configure X11 Forwarding Using SSH In Linux

NEW QUESTION 200
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 201
After connecting to a remote host via SSH, an administrator attempts to run an application but receives the following error:
[user@workstation ~]$ ssh admin@srv1 Last login: Tue Mar 29 18:03:34 2022
[admin@srvl ~] $ /usr/local/bin/config_manager Error: cannot open display:
[admin@srv1 ~] $
Which of the following should the administrator do to resolve this error?

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

Answer: A

Explanation:
The error indicates that the application requires an X11 display, but the SSH session does not forward the X11 connection. To enable X11 forwarding, the administrator needs to use the ssh -X option, which requests X11 forwarding with authentication spoofing. This will set the DISPLAY environment variable on the remote host and allow the application to open a window on the local display.
References
? CompTIA Linux+ (XK0-005) Certification Study Guide, page 314
? Open a window on a remote X display (why "Cannot open display")?, answer by Gilles 'SO- stop being evil'

NEW QUESTION 204
A systems administrator detected corruption in the /data filesystem. Given the following output:

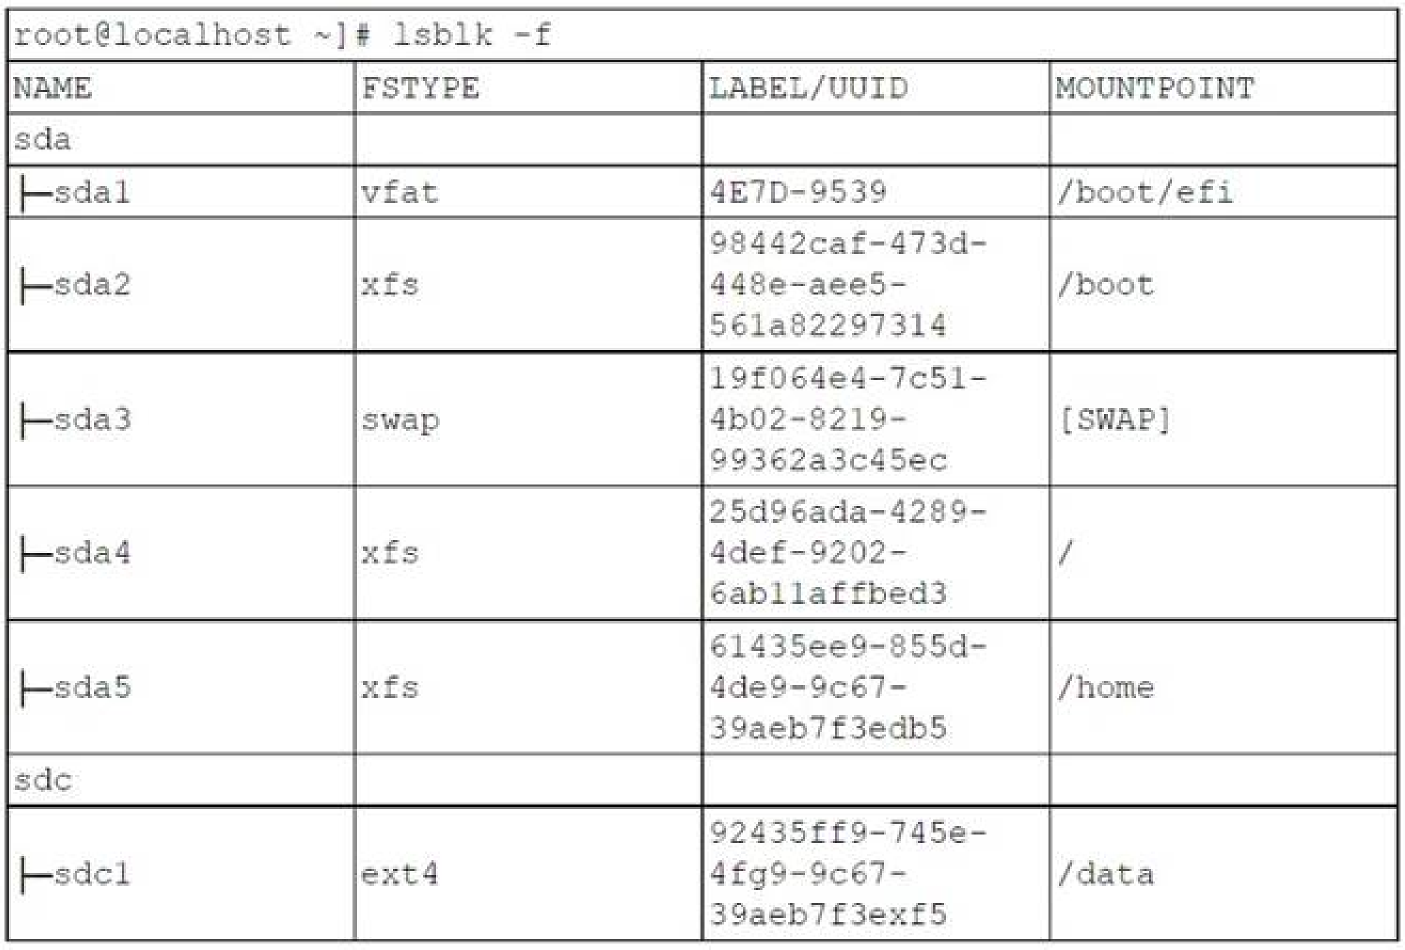

root@localhost ~]# lsblk -f

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

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

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

Answer: B

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

NEW QUESTION 208
A Linux administrator has logged in to a server for the first time and needs to know which services are allowed through the firewall. Which of the following options will return the results for which the administrator is looking?

A. firewall-cmd —get-services
B. firewall-cmd —check-config
C. firewall-cmd —list-services
D. systemct1 status firewalld

Answer: C

Explanation:
The firewall-cmd --list-services command will return the results for which the administrator is looking. This command will list all services that are allowed through the firewall in the default zone or a specified zone. A service is a predefined set of ports and protocols that can be enabled or disabled by firewalld. The firewall-cmd --get-services command will list all available services that are supported by firewalld, not only those that are allowed through the firewall. The firewall-cmd --check-config command will check if firewalld configuration files are valid, not list services. The systemct1 status firewalld command will display information about the firewalld service unit, such as its state, PID, memory usage, and logs, not list services. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 543.

NEW QUESTION 211
A Linux administrator would like to use systemd to schedule a job to run every two hours. The administrator creates timer and service definitions and restarts the server to load these new configurations. After the restart, the administrator checks the log file and notices that the job is only running daily. Which of the following is MOST likely causing the issue?

A. The checkdiskspace.service is not running.
B. The checkdiskspace.service needs to be enabled.
C. The OnCalendar schedule is incorrect in the timer definition.
D. The system-daemon services need to be reloaded.

Answer: C

Explanation:
The OnCalendar schedule is incorrect in the timer definition, which is causing the issue. The OnCalendar schedule defines when the timer should trigger the service. The format of the schedule is OnCalendar=<year>-<month>-<day> <hour>:<minute>:<second>. If any of the fields are omitted, they are assumed to be *,

which means any value. Therefore, the schedule OnCalendar=*-*-* 00:00:00 means every day at midnight, which is why the job is running daily. To make the job run every two hours, the schedule should be OnCalendar=*-*-* *:00:00/2, which means every hour divisible by 2 at the start of the minute. The other options are incorrect because they are not related to the schedule. The checkdiskspace.service is running, as shown by the output of systemct1 status checkdiskspace.service. The checkdiskspace.service is enabled, as shown by the output of systemct1 is-enabled checkdiskspace.service. The system-daemon services do not need to be reloaded, as the timer and service definitions are already loaded by the
restart. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, page 437.

NEW QUESTION 214

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

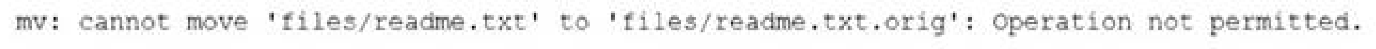

```
mv: cannot move 'files/readme.txt' to 'files/readme.txt.orig': Operation not permitted.
```

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

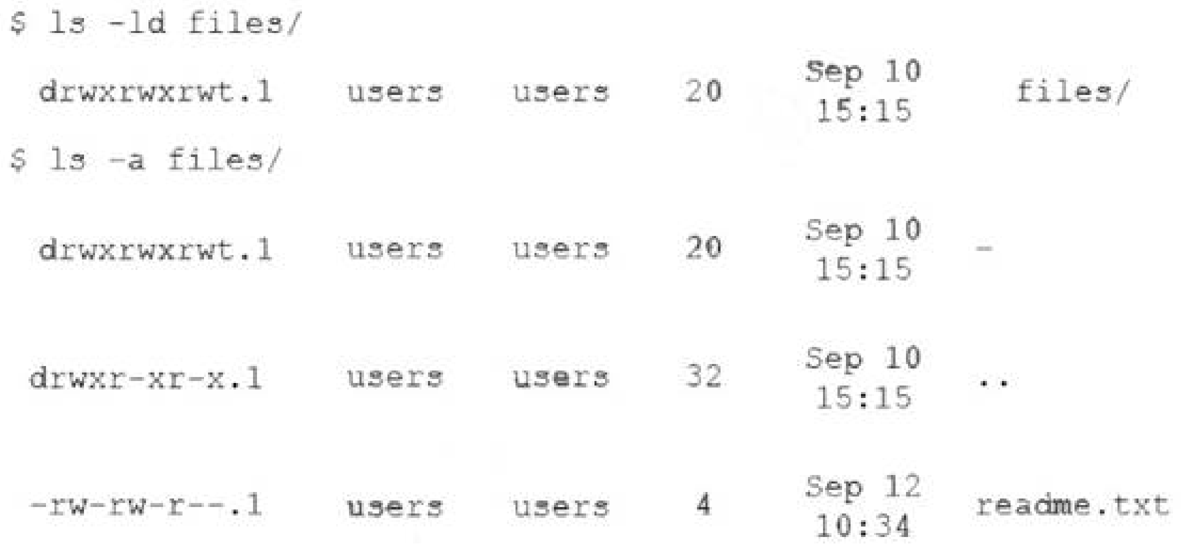

```
$ ls -ld files/
  drwxrwxrwt.1    users    users    20    Sep 10       files/
                                          15:15
$ ls -a files/
  drwxrwxrwt.1    users    users    20    Sep 10    -
                                          15:15

 drwxr-xr-x.1     users    users    32    Sep 10    ..
                                          15:15

 -rw-rw-r--.1     users    users    4     Sep 12    readme.txt
                                          10:34
```

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 217

......