XK0-005 Dumps

CompTIA Linux+ Certification Exam

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

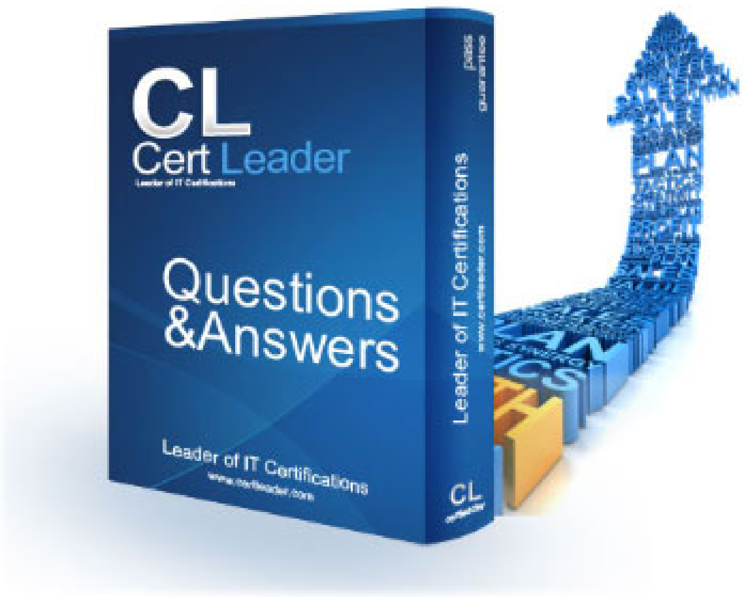

NEW QUESTION 1
A Linux administrator was notified that a virtual server has an I/O bottleneck. The Linux administrator analyzes the following output:

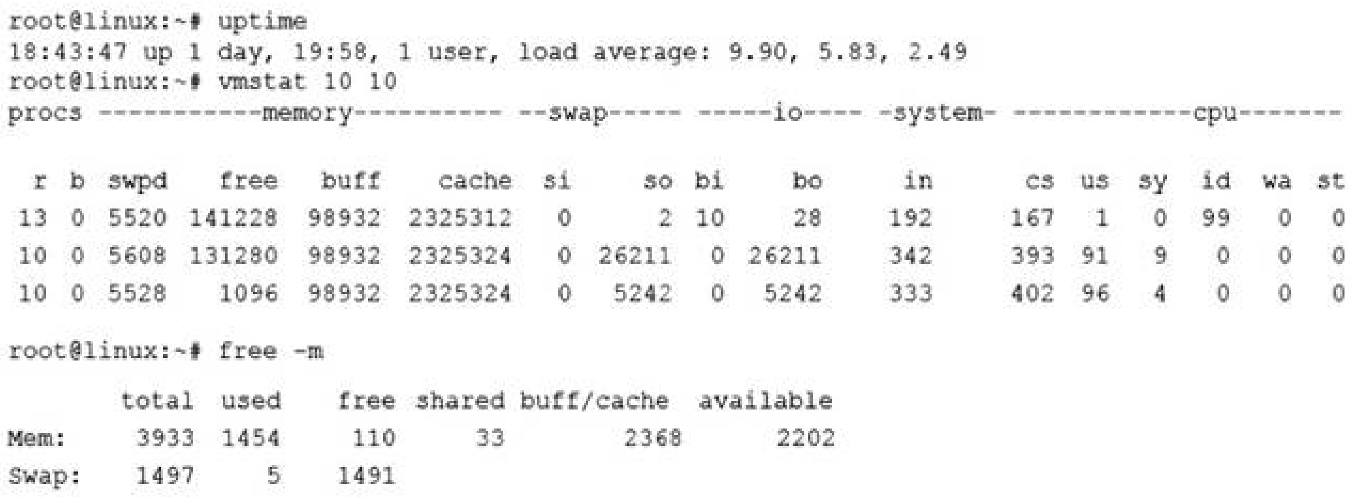

```
root@linux:~# uptime
18:43:47 up 1 day, 19:58, 1 user, load average: 9.90, 5.83, 2.49
root@linux:~# vmstat 10 10
procs -----------memory---------- --swap----- -----io---- -system- ------------cpu-------

  r  b  swpd    free   buff    cache  si      so  bi     bo      in       cs  us  sy  id  wa  st
 13  0  5520  141228  98932  2325312   0       2  10     28     192      167   1   0  99   0   0
 10  0  5608  131280  98932  2325324   0   26211   0  26211     342      393  91   9   0   0   0
 10  0  5528    1096  98932  2325324   0    5242   0   5242     333      402  96   4   0   0   0

root@linux:~# free -m

        total  used    free  shared buff/cache  available
Mem:     3933  1454     110      33       2368       2202
Swap:    1497     5    1491
```

Given there is a single CPU in the sever, which of the following is causing the slowness?

A. The system is running out of swap space.
B. The CPU is overloaded.
C. The memory is exhausted.
D. The processes are paging.

Answer: B

Explanation:
The slowness is caused by the CPU being overloaded. The iostat command shows that the CPU utilization is 100%, which means that there are more processes competing for CPU time than the CPU can handle. The other options are incorrect because:
? The system is not running out of swap space, as shown by the iostat command, which shows that there is no swap activity (si and so columns are zero).
? The memory is not exhausted, as shown by the free -m command, which shows that there is still available memory (avail column) and free buffer/cache memory (buff/cache column).
? The processes are not paging, as shown by the vmstat command, which shows that there are no major page faults (majflt column) and no swap activity (si and so columns). References: CompTIA Linux+ Study Guide, Fourth Edition, page 417- 419, 424-425.

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

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

Answer: D

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

NEW QUESTION 4
A Linux administrator is tasked with creating resources using containerization. When deciding how to create this type of deployment, the administrator identifies some key features, including portability, high availability, and scalability in production. Which of the following should the Linux administrator choose for the new

design?

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

Answer: D

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

NEW QUESTION 5
A Linux administrator is troubleshooting a systemd mount unit file that is not working correctly. The file contains:
[root@system] # cat mydocs.mount [Unit]
Description=Mount point for My Documents drive [Mount]
What=/dev/drv/disk/by-uuid/94afc9b2-ac34-ccff-88ae-297ab3c7ff34 Where=/home/user1/My Documents
Options=defaults Type=xfs
[Install]
WantedBy=multi-user.target
The administrator verifies the drive UUID correct, and user1 confirms the drive should be mounted as My Documents in the home directory. Which of the following can the administrator do to fix the issues with mounting the drive? (Select two).

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

Answer: AE

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

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

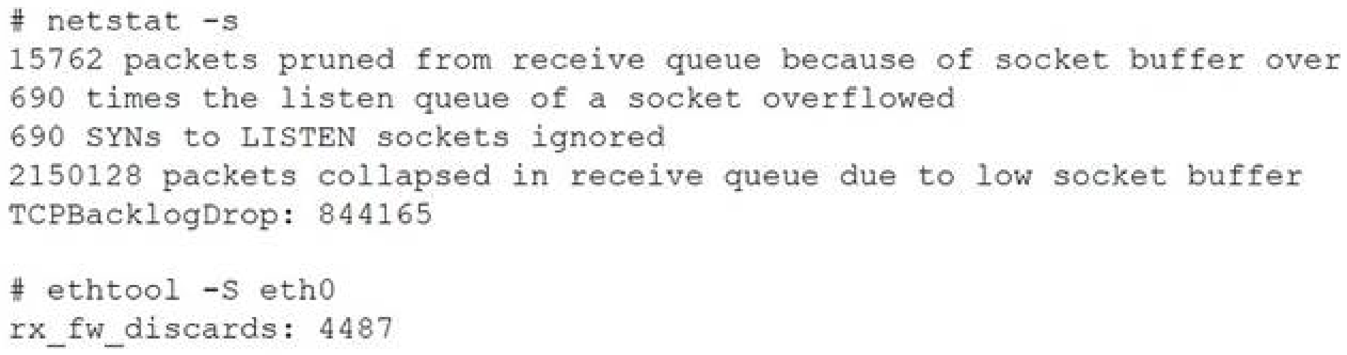

```
# netstat -s
15762 packets pruned from receive queue because of socket buffer over
690 times the listen queue of a socket overflowed
690 SYNs to LISTEN sockets ignored
2150128 packets collapsed in receive queue due to low socket buffer
TCPBacklogDrop: 844165

# ethtool -S eth0
rx_fw_discards: 4487
```

Which of the following commands will improve the latency issue?

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

Answer: D

Explanation:
The best command to use to improve the latency issue is D. # echo ‘net.core.rmem max = 12500000’ >> /etc/sysctl.conf # echo ‘net.core.wmem_max = 12500000’ >> /etc/sysctl.conf # sysctl -p. This command will increase the size of the receive and send buffers for the network interface, which can improve the network performance and reduce packet loss. The sysctl command will apply the changes to the kernel parameters without rebooting the system.
The other commands are either incorrect or not suitable for this task. For example:
? A. # echo ‘net.core.net_backlog = 5000000’ >> /etc/sysctl.conf # sysctl -p # systemctl daemon-reload will try to increase the backlog queue for incoming connections, but this is not relevant for the latency issue. The systemctl daemon- reload command is also unnecessary, as it only reloads the systemd configuration files, not the kernel parameters.
? B. # ifdown eth0 # ip link set dev eth0 mtu 800 # ifup eth0 will try to change the maximum transmission unit (MTU) of the network interface to 800 bytes, but this is too low and may cause fragmentation and performance degradation. The default MTU for Ethernet is 1500 bytes, and it should not be changed unless there is a specific reason.
? C. # systemctl stop network # ethtool -g eth0 512 # systemctl start network will try to change the ring buffer size of the network interface to 512, but this is too small and may cause packet drops and latency spikes. The default ring buffer size for Ethernet is usually 4096 or higher, and it should be increased if there is a high network traffic.

NEW QUESTION 7
A cloud engineer is installing packages during VM provisioning. Which of the following should the engineer use to accomplish this task?

A. Cloud-init
B. Bash
C. Docker
D. Sidecar

Answer: A

Explanation:
The cloud engineer should use cloud-init to install packages during VM provisioning. Cloud-init is a tool that allows the customization of cloud instances at boot time. Cloud-init can perform various tasks, such as setting the hostname, creating users, installing packages, configuring network, and running scripts. Cloud-init can work with different cloud platforms and Linux distributions. This is the correct tool to accomplish the task. The other options are incorrect because they are either not suitable for cloud provisioning (Bash or Docker) or not a tool but a design pattern (Sidecar). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 563.

NEW QUESTION 8
An administrator needs to make an application change via a script that must be run only in console mode. Which of the following best represents the sequence the administrator should execute to accomplish this task?

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

Answer: A

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

NEW QUESTION 9
A systems engineer has deployed a new application server, but the server cannot communicate with the backend database hostname. The engineer confirms that the application server can ping the database server's IP address. Which of the following is the most likely cause of the issue?

A. Incorrect DNS servers
B. Unreachable default gateway
C. Missing route configuration
D. Misconfigured subnet mask

Answer: A

Explanation:
This is because the application server can ping the database server’s IP address, but not its hostname, which suggests that the DNS resolution is not working properly. DNS servers are responsible for translating hostnames into IP addresses, and vice versa. If the application server has incorrect or unreachable DNS servers configured, it will not be able to resolve the hostname of the database server and communicate with it.
To troubleshoot this issue, the systems engineer should check the DNS configuration on the application server, which is usually stored in the /etc/resolv.conf file. This file should contain valid nameserver entries that point to the DNS servers that can resolve the database server’s hostname. For example, a typical /etc/resolv.conf file may look like this: nameserver 8.8.8.8 nameserver 8.8.4.4
These are the IP addresses of Google’s public DNS servers, which can be used as a fallback option if the default DNS servers are not working.
Alternatively, the systems engineer can use the nslookup or dig commands to test the DNS resolution of the database server’s hostname from the application server. These commands will query a specified DNS server and return the IP address of the hostname, or an error message if the resolution fails. For example, to query Google’s public DNS server for the IP address of comptia.org, the command would be:
nslookup comptia.org 8.8.8.8 or dig comptia.org @8.8.8.8

NEW QUESTION 10

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

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

Answer: B

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

NEW QUESTION 13
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 18
A systems administrator wants to be sure the sudo rules just added to /etc/sudoers are valid. Which of the following commands can be used for this task?

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

Answer: A

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

NEW QUESTION 20
A systems administrator is gathering information about a file type and the contents of a file. Which of the following commands should the administrator use to

accomplish this task?

A. file filename
B. touch filename
C. grep filename
D. lsof filename

Answer: A

Explanation:
The file command is used to determine the type of a file by examining its contents. It can recognize many different formats, such as text, binary, executable, compressed, image, audio, video, etc. It can also display some additional information about the file, such as encoding, size, dimensions, etc12
References: 1: file(1) - Linux manual page 2: How to use the file command in Linux

NEW QUESTION 25
Users report that connections to a MariaDB service are being closed unexpectedly. A systems administrator troubleshoots the issue and finds the following message in /var/log/messages:

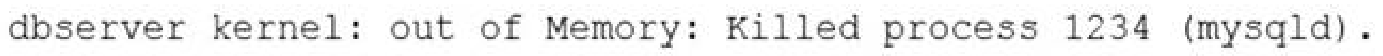

```
dbserver kernel: out of Memory: Killed process 1234 (mysqld).
```

Which of the following is causing the connection issue?

A. The process mysqld is using too many semaphores.
B. The server is running out of file descriptors.
C. Something is starving the server resources.
D. The amount of RAM allocated to the server is too high.

Answer: B

Explanation:
The message in /var/log/messages indicates that the server is running out of file descriptors. A file descriptor is a non-negative integer identifier for an open file in Linux. Each process has a table of open file descriptors where a new entry is appended upon opening a new file. There is a limit on how many file descriptors a process can open at a time, which depends on the system configuration and the user privileges. If a process tries to open more files than the limit, it will fail with an error message like "Too many open files". This could cause connections to be closed unexpectedly or other problems with the application.
The other options are not correct causes for the connection issue. The process mysqld is not using too many semaphores, which are synchronization mechanisms for processes that share resources. Semaphores are not related to file descriptors or open files. Something is not starving the server resources, which could mean high CPU usage, memory pressure, disk I/O, network congestion, or other factors that affect performance. These could cause slowdowns or timeouts, but not file descriptor exhaustion. The amount of RAM allocated to the server is not too high, which could cause swapping or paging if it exceeds the physical memory available. This could also affect performance, but not file descriptor availability. References: File Descriptor Requirements (Linux Systems); Limits on the Number of Linux File Descriptors

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

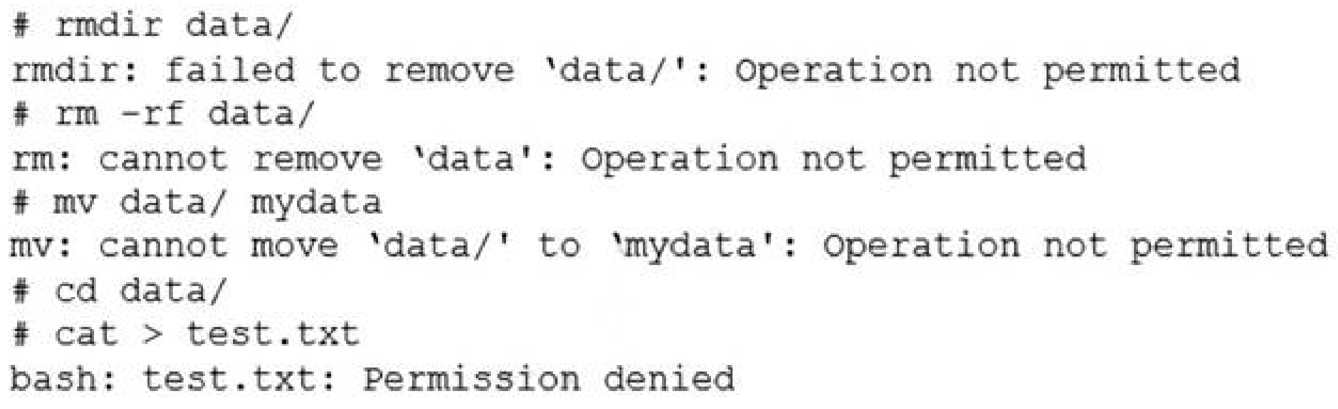

```
# rmdir data/
rmdir: failed to remove 'data/': Operation not permitted
# rm -rf data/
rm: cannot remove 'data': Operation not permitted
# mv data/ mydata
mv: cannot move 'data/' to 'mydata': Operation not permitted
# cd data/
# cat > test.txt
bash: test.txt: Permission denied
```

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

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

Answer: C

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

NEW QUESTION 31

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

Output 1:

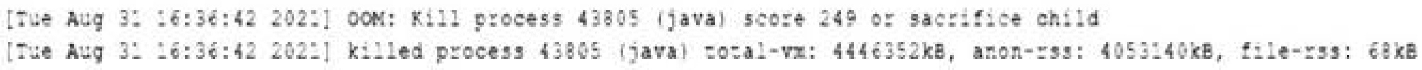

```
[Tue Aug 31 16:36:42 2021] OOM: Kill process 43805 (java) score 249 or sacrifice child
[Tue Aug 31 16:36:42 2021] killed process 43805 (java) total-vm: 4446352kB, anon-rss: 4053140kB, file-rss: 68kB
```

Output 2:

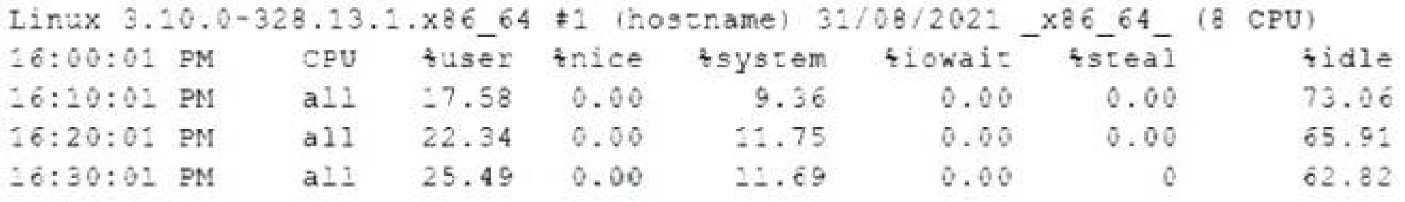

```
Linux 3.10.0-328.13.1.x86_64 #1 (hostname) 31/08/2021 _x86_64_ (8 CPU)
16:00:01 PM     CPU    %user   %nice   %system   %iowait   %steal      %idle
16:10:01 PM     all    17.58   0.00       9.36      0.00     0.00      73.06
16:20:01 PM     all    22.34   0.00      11.75      0.00     0.00      65.91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

Output 3:

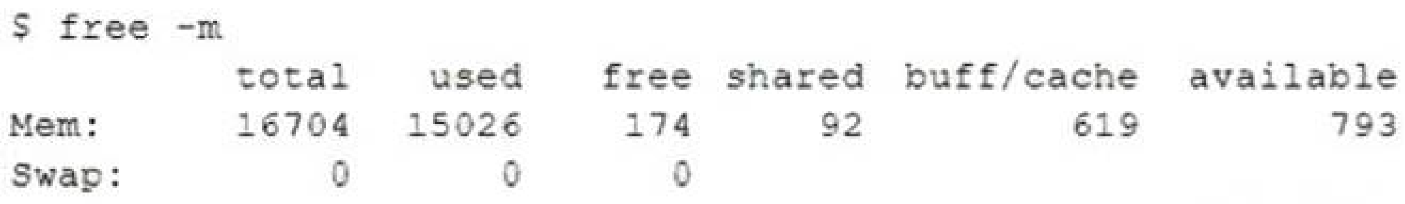

```
$ free -m
           total    used    free  shared  buff/cache  available
Mem:       16704   15026     174      92         619        793
Swap:          0       0       0
```

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

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

Answer: D

Explanation:
Based on the output of the image sent by the user, the system requires more swap space to allow for spikes in usage during peak hours and prevent the Java process from stopping due to a lack of memory. The output shows that there is only 0 MB of swap space available on the system, which means that there is no room for swapping out memory pages when physical memory is full or low. The output also shows that there is only 793 MB of available memory on the system, which may not be enough to handle high- demand applications such as Java. This may cause Java to stop working due to insufficient memory or trigger an OutOfMemoryError exception. Configuring more swap space on the system would help to alleviate this issue by providing more virtual memory for applications and improving performance. Configuring memory allocation policies during business hours will not help to solve this issue, as it will not increase the amount of available memory or swap space on the system. Configuring a different nice value for Java process will not help to solve this issue, as it will only affect its scheduling priority, not its memory consumption or allocation. Configuring more CPU cores will not help to solve this issue, as it will only increase processing power, not memory capacity or availability. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Memory and Process Execution, page 468.

NEW QUESTION 35
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 40
A Linux system fails to start and delivers the following error message:

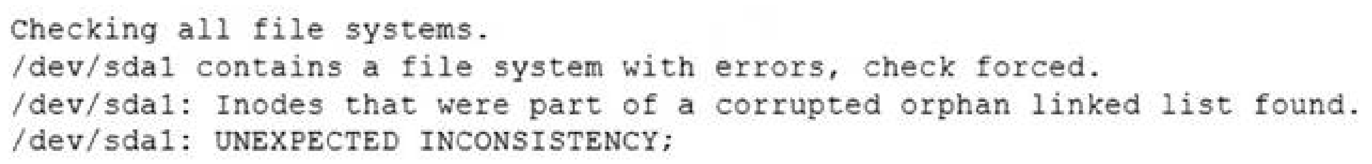

```
Checking all file systems.
/dev/sda1 contains a file system with errors, check forced.
/dev/sda1: Inodes that were part of a corrupted orphan linked list found.
/dev/sda1: UNEXPECTED INCONSISTENCY;
```

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

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

Answer: A

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

NEW QUESTION 44
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 49
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:

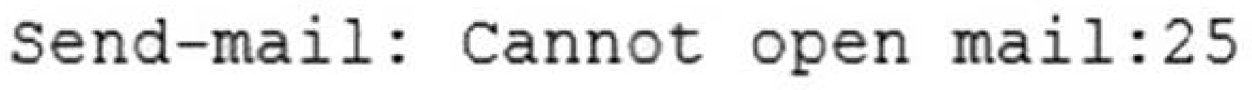

```
Send-mail: Cannot open mail:25
```

The local machine DNS settings are:

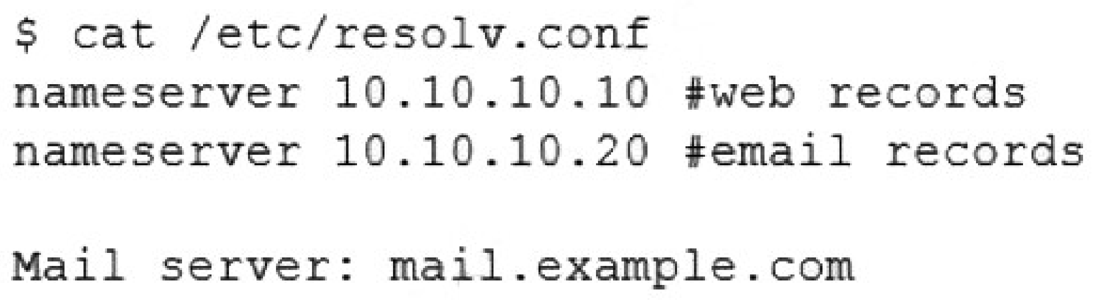

```
$ cat /etc/resolv.conf
nameserver 10.10.10.10 #web records
nameserver 10.10.10.20 #email records

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 53
Users have been unable to save documents to /home/tmp/temp and have been receiving the following error:
Path not found
A junior technician checks the locations and sees that /home/tmp/tempa was accidentally created instead of /home/tmp/temp. Which of the following commands should the technician use to fix this issue?

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

Answer: B

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

NEW QUESTION 57
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 58
Which of the following is the best tool for dynamic tuning of kernel parameters?

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

Answer: A

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

NEW QUESTION 63
A Linux administrator is trying to remove the ACL from the file /home/user/data. txt but receives the following error message:

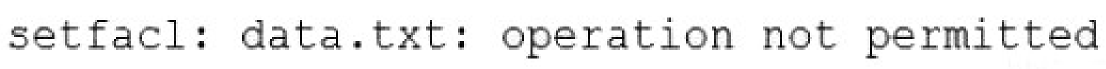

```
setfacl: data.txt: operation not permitted
```

Given the following analysis:

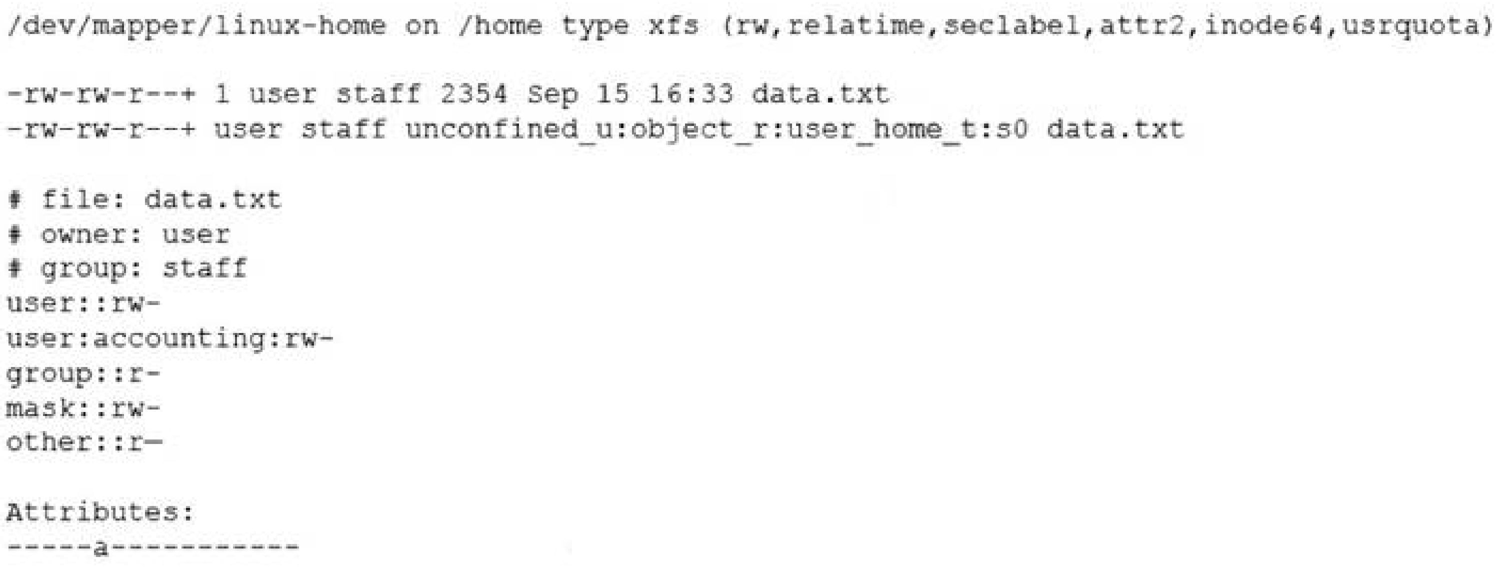

```
/dev/mapper/linux-home on /home type xfs (rw,relatime,seclabel,attr2,inode64,usrquota)

-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 64
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 69
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:

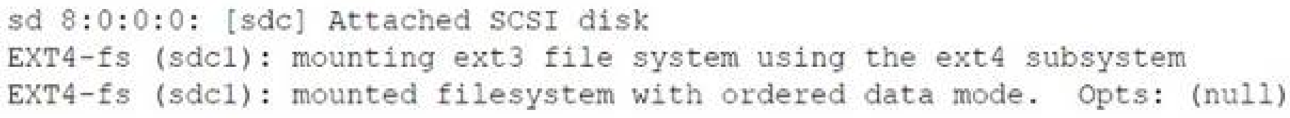

```
sd 8:0:0:0: [sdc] Attached SCSI disk
EXT4-fs (sdcl): mounting ext3 file system using the ext4 subsystem
EXT4-fs (sdcl): mounted filesystem with ordered data mode.  Opts: (null)
```

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 74
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.

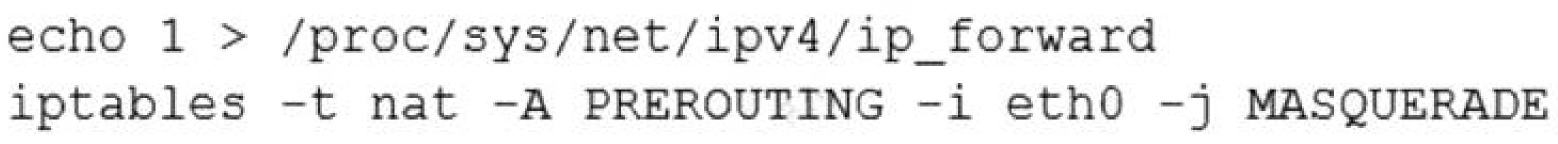

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A PREROUTING -i eth0 -j MASQUERADE
```

A.

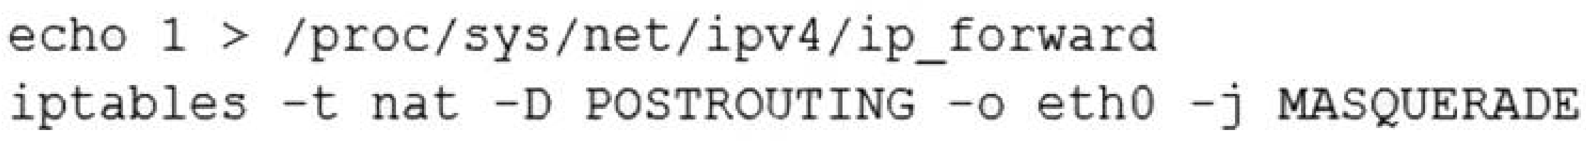

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -D POSTROUTING -o eth0 -j MASQUERADE
```

B.

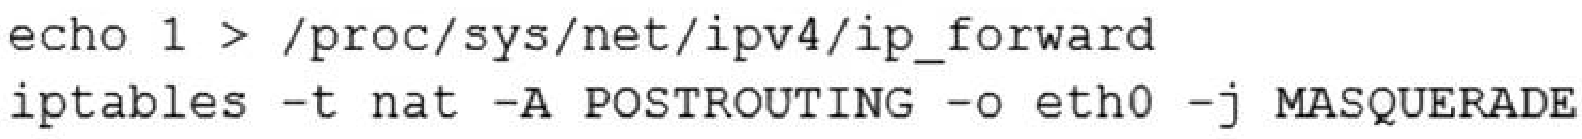

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
```

C.

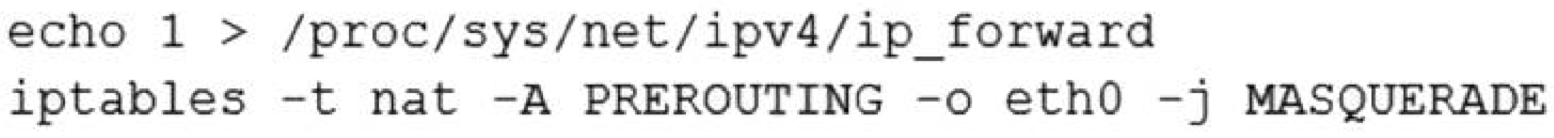

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A PREROUTING -o eth0 -j MASQUERADE
```

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 77
A systems administrator is tasked with preventing logins from accounts other than root, while the file /etc/nologin exists. Which of the following PAM modules will accomplish this task?

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

Answer: D

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

NEW QUESTION 80
A systems administrator checked out the code from the repository, created a new branch, made changes to the code, and then updated the main branch. The systems administrator wants to ensure that the Terraform state files do not appear in the main branch. Which of following should the administrator use to meet this requirement?

A. clone
B. gitxgnore
C. get
D. .ssh

Answer: B

Explanation:
To prevent certain files from being tracked by Git, the administrator can use a .gitignore file (B) in the repository. The .gitignore file can specify patterns of files or directories that Git should ignore. This way, the Terraform state files will not appear in the main branch or any other branch. The other commands are not related to this requirement. References:
? [CompTIA Linux+ Study Guide], Chapter 10: Working with Git, Section: Ignoring Files with .gitignore
? [How to Use .gitignore File]

NEW QUESTION 81
A systems administrator is receiving tickets from users who cannot reach the application app that should be listening on port 9443/tcp on a Linux server.
To troubleshoot the issue, the systems administrator runs netstat and receives the following output:

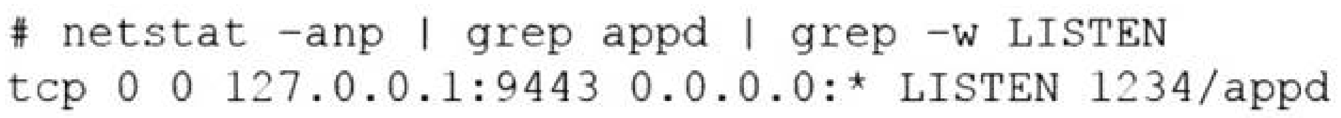

```
# netstat -anp | grep appd | grep -w LISTEN
tcp 0 0 127.0.0.1:9443 0.0.0.0:* LISTEN 1234/appd
```

Based on the information above, which of the following is causing the issue?

A. The IP address 0.0.0.0 is not valid.
B. The application is listening on the loopback interface.
C. The application is listening on port 1234.
D. The application is not running.

Answer: B

Explanation:
The server is in a "Listen" state on port 9943 using its loopback address. The "1234" is a process-id
The cause of the issue is that the application is listening on the loopback interface. The loopback interface is a virtual network interface that is used for internal communication within the system. The loopback interface has the IP address 127.0.0.1, which is also known as localhost. The netstat output shows that the application is listening on port 9443 using the IP address 127.0.0.1. This means that the application can only accept connections from the same system, not from other systems on the network. This can prevent the users from reaching the application and cause the issue. The administrator should configure the application to listen on the IP address 0.0.0.0, which means all available interfaces, or on the specific IP address of the system that is reachable from the network. This will allow the application to accept connections from other systems and resolve the issue. The cause of the issue is that the application is listening on the loopback interface. This is the correct answer to the question. The other options are incorrect because they are not supported by the outputs. The IP address 0.0.0.0 is valid and means all interfaces, the application is not listening on port 1234, and the application is running as shown by the process ID 1234. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 383.

NEW QUESTION 83
A Linux administrator cloned an existing Linux server and built a new server from that clone. The administrator encountered the following error after booting the cloned server:

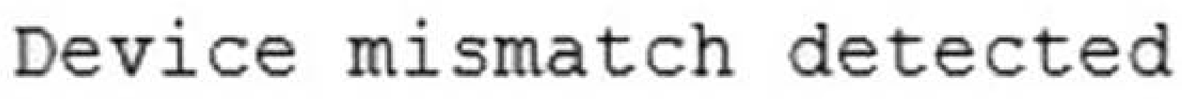

```
Device mismatch detected
```

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

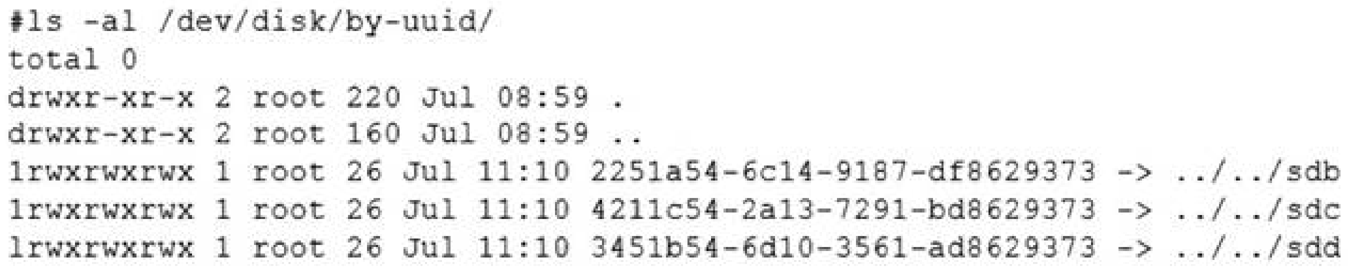

```
#ls -al /dev/disk/by-uuid/
total 0
drwxr-xr-x 2 root 220 Jul 08:59 .
drwxr-xr-x 2 root 160 Jul 08:59 ..
lrwxrwxrwx 1 root 26 Jul 11:10 2251a54-6c14-9187-df8629373 -> ../../sdb
lrwxrwxrwx 1 root 26 Jul 11:10 4211c54-2a13-7291-bd8629373 -> ../../sdc
lrwxrwxrwx 1 root 26 Jul 11:10 3451b54-6d10-3561-ad8629373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

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

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

Answer: A

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

NEW QUESTION 86
An administrator is trying to diagnose a performance issue and is reviewing the following output:

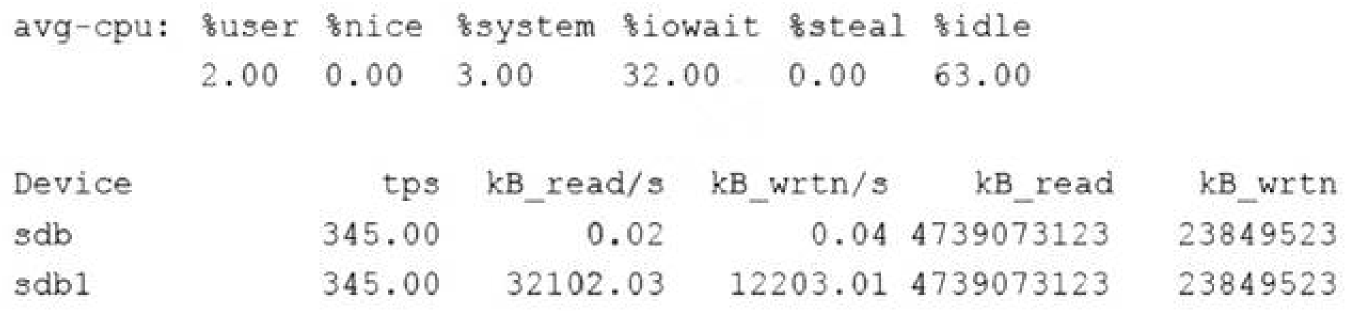

```
avg-cpu: %user %nice %system %iowait %steal %idle
          2.00  0.00   3.00   32.00   0.00  63.00

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 88
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 89
A junior systems administrator recently installed an HBA card in one of the servers that is deployed for a production environment. Which of the following commands can the administrator use to confirm on which server the card was installed?

A. lspci | egrep 'hba| fibr'
B. lspci | zgrep 'hba | fibr'
C. lspci | pgrep 'hba| fibr'
D. lspci | 'hba | fibr'

Answer: A

Explanation:
The best command to use to confirm on which server the HBA card was installed is A. lspci | egrep ‘hba| fibr’. This command will list all the PCI devices on the server and filter the output for those that match the pattern ‘hba’ or ‘fibr’, which are likely to be related to the HBA card. The egrep command is a variant of grep that supports extended regular expressions, which allow the use of the ‘|’ operator for alternation. The other commands are either invalid or will not produce the desired output. For example:
? B. lspci | zgrep ‘hba | fibr’ will try to use zgrep, which is a command for searching compressed files, not standard output.
? C. lspci | pgrep ‘hba| fibr’ will try to use pgrep, which is a command for finding processes by name or other attributes, not text patterns.
? D. lspci | ‘hba | fibr’ will try to use ‘hba | fibr’ as a command, which is not valid and will cause an error.

NEW QUESTION 92
A systems administrator is tasked with mounting a USB drive on a system. The USB drive has a single partition, and it has been mapped by the system to the device /dev/sdb. Which of the following commands will mount the USB to /media/usb?

A. mount /dev/sdb1 /media/usb
B. mount /dev/sdb0 /media/usb

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

Answer: A

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

NEW QUESTION 96
A Linux administrator needs to transfer a local file named accounts . pdf to a remote / tmp directory of a server with the IP address 10.10.10.80. Which of the following commands needs to be executed to transfer this file?

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

Answer: B

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

NEW QUESTION 99
A user created the following script file:
# ! /bin/bash
# FILENAME: /home/user/ script . sh echo "hello world"
exit 1
However, when the user tried to run the script file using the command "script . sh, an error returned indicating permission was denied. Which of the follow-ing should the user execute in order for the script to run properly?

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

Answer: A

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

NEW QUESTION 104
The development team created a new branch with code changes that a Linux administrator needs to pull from the remote repository. When the administrator looks for the branch in Git, the branch in question is not visible. Which of the following commands should the Linux administrator run to refresh the branch information?

A. git fetch
B. git checkout
C. git clone
D. git branch

Answer: A

Explanation:
The git fetch command downloads commits, files, and refs from a remote repository into the local one. It also updates the remote-tracking branches, which are references to the state of the remote branches. By running git fetch, the administrator can see the new branch created by the development team and then use git checkout to switch to it12. References: 1: Git - git-fetch Documentation 2: Git Fetch | Atlassian Git Tutorial

NEW QUESTION 107
A systems administrator is encountering performance issues. The administrator runs 3 commands with the following output

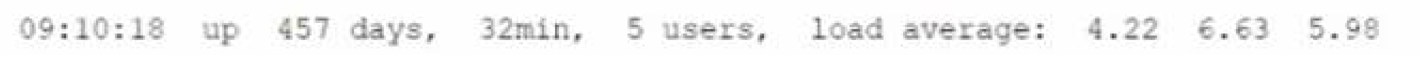

```
09:10:18 up 457 days, 32min, 5 users, load average: 4.22 6.63 5.98
```

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

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

Answer: A

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

NEW QUESTION 109
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 112
An administrator needs to increase the system priority of a process with PID 2274. Which of the following commands should the administrator use to accomplish this task?

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

Answer: A

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

NEW QUESTION 114
A Linux administrator needs to redirect all HTTP traffic temporarily to the new proxy server 192.0.2.25 on port 3128. Which of the following commands will accomplish this task?

A. iptables -t nat -D PREROUTING -p tcp --sport 80 -j DNAT - -to-destination 192.0.2.25:3128
B. iptables -t nat -A PREROUTING -p top --dport 81 -j DNAT –-to-destination 192.0.2.25:3129
C. iptables -t nat -I PREROUTING -p top --sport 80 -j DNAT –-to-destination 192.0.2.25:3129
D. iptables -t nat -A PREROUTING -p tcp --dport 80 -j DNAT –-to-destination 192.0.2.25:3128

Answer: D

Explanation:
The command iptables -t nat -A PREROUTING -p tcp --dport 80 -j DNAT -- to-destination 192.0.2.25:3128 adds a rule to the nat table that redirects all incoming TCP packets with destination port 80 (HTTP) to the proxy server 192.0.2.25 on port 3128. This is the correct way to achieve the task. The other options are incorrect because they either delete a rule (-D), use the wrong protocol (top instead of tcp), or use the wrong port (81 instead of 80). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 381.

NEW QUESTION 118
A Linux system is failing to boot with the following error:

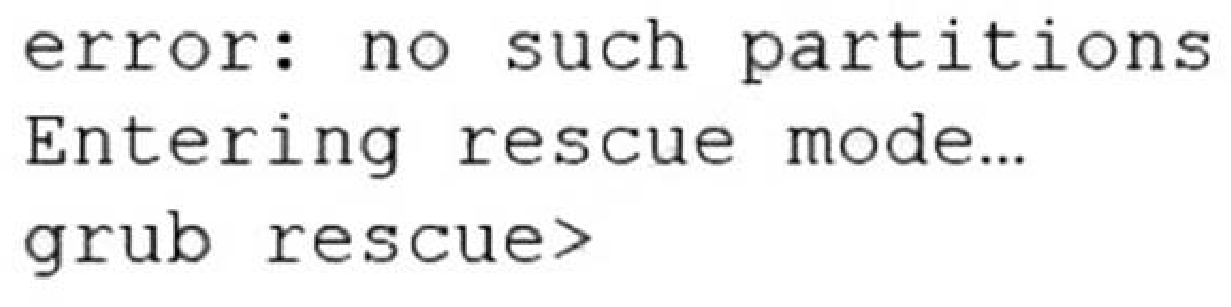

```
error: no such partitions
Entering rescue mode...
grub rescue>
```

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 121
A cloud engineer is asked to copy the file deployment.yaml from a container to the host where the container is running. Which of the following commands can accomplish this task?

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

Answer: B

Explanation:
The command docker cp container_id:/deployment.yaml deployment.yaml can accomplish the task of copying the file deployment.yaml from a container to the host.
The docker command is a tool for managing Docker containers and images. The cp option copies files or directories between a container and the local filesystem. The container_id is the identifier of the container, which can be obtained by using the docker ps command.
The /deployment.yaml is the path of the file in the container, which must be preceded by a slash. The deployment.yaml is the path of the file on the host, which can be relative or absolute. The command docker cp container_id:/deployment.yaml deployment.yaml will copy the file deployment.yaml from the container to the current working directory on the host. This is the correct command to use to accomplish the task. The other options are incorrect because they either use the wrong syntax (docker cp container_id/deployment.yaml deployment.yaml or docker cp container_id/deployment.yaml local://deployment.yaml) or do not exist (docker cp deployment.yaml local://deployment.yaml). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 567.

NEW QUESTION 126
A systems administrator creates a public key for authentication. Which of the following tools is most suitable to use when uploading the key to the remote servers?

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

Answer: B

Explanation:
The best tool to use when uploading the public key to the remote servers is
* B. ssh-copy-id. This tool will copy the public key from the local computer to the remote server and append it to the authorized_keys file, which is used for public

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

NEW QUESTION 129
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 131
An administrator accidentally deleted the /boot/vmlinuz file and must resolve the issue before the server is rebooted. Which of the following commands should the administrator use to identify the correct
version of this file?

A. rpm -qa | grep kernel; uname -a
B. yum -y update; shutdown -r now
C. cat /etc/centos-release; rpm -Uvh --nodeps
D. telinit 1; restorecon -Rv /boot

Answer: A

Explanation:
The command rpm -qa | grep kernel lists all the installed kernel packages, and the command uname -a displays the current kernel version. These commands can help the administrator identify the correct version of the /boot/vmlinuz file, which is the kernel image file. The other options are not relevant or helpful for this task.
References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing the Linux Boot Process, page 267.

NEW QUESTION 135
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 138
Which of the following files holds the system configuration for journal when running systemd?

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

Answer: A

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

NEW QUESTION 143
A systems administrator is tasked with changing the default shell of a system account in order to disable iterative logins. Which of the following is the best option for the administrator to use as the new shell?

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

Answer: A

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

NEW QUESTION 146
A systems administrator is installing various software packages using a pack-age manager. Which of the following commands would the administrator use on the Linux server to install the package?

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

Answer: D

NEW QUESTION 149
After listing the properties of a system account, a systems administrator wants to remove the expiration date of a user account. Which of the following commands will accomplish this task?

A. chgrp system accountname
B. passwd –s accountname
C. chmod -G system account name
D. chage -E -1 accountname

Answer: D

Explanation:
The command chage -E -1 accountname will accomplish the task of removing the expiration date of a user account. The chage command is a tool for changing user password aging information on Linux systems. The -E option sets the expiration date of the user account, and the -1 value means that the account will never expire. The command chage -E -1 accountname will remove the expiration date of the user account named accountname. This is the correct command to use to accomplish the task. The
other options are incorrect because they either do not affect the expiration date
(chgrp, passwd, or chmod) or do not exist (chmod -G). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Users and Groups, page 467.

NEW QUESTION 151
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:

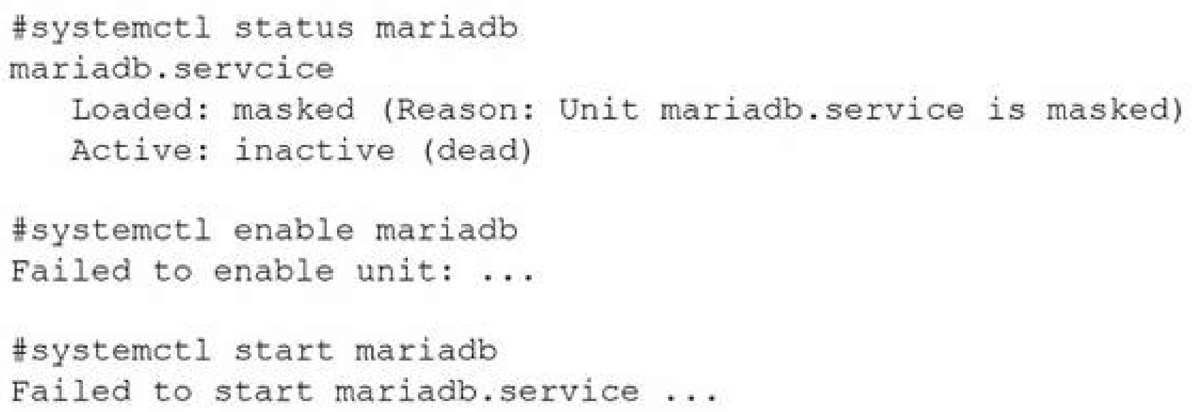

```
#systemctl status mariadb
mariadb.servcice
   Loaded: masked (Reason: Unit mariadb.service is masked)
   Active: inactive (dead)

#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 154
A Linux system is failing to start due to issues with several critical system processes. Which of the following options can be used to boot the system into the single user mode? (Choose two.)

A. Execute the following command from the GRUB rescue shell: mount -o remount, ro/sysroot.
B. Interrupt the boot process in the GRUB menu and add systemd.unit=single in the kernel line.
C. Interrupt the boot process in the GRUB menu and add systemd.unit=rescue.target in the kernel line.
D. Interrupt the boot process in the GRUB menu and add single=user in the kernel line.
E. Interrupt the boot process in the GRUB menu and add init=/bin/bash in the kernel line.
F. Interrupt the boot process in the GRUB menu and add systemd.unit=single.target in the kernel line.

Answer: CF

Explanation:
The administrator can use the following two options to boot the system into the single user mode:
? Interrupt the boot process in the GRUB menu and add systemd.unit=rescue.target in the kernel line. This option will boot the system into the rescue mode, which is a minimal environment that allows the administrator to perform basic tasks such as repairing the system. The GRUB menu is a screen that appears when the system is powered on and allows the administrator to choose which kernel or operating system to boot. The kernel line is a line that specifies the parameters for the kernel, such as the root device, the init system, and the boot options. The administrator can interrupt the boot process by pressing the e key in the GRUB menu and edit the kernel line by adding systemd.unit=rescue.target at the end. This option will tell the system to use the rescue target, which is a unit that defines the state of the system in the rescue mode. The administrator can then press Ctrl+X to boot the system with the modified kernel line. This option will boot the system into the single user mode and allow the administrator to troubleshoot the issues.
? Interrupt the boot process in the GRUB menu and add systemd.unit=single.target in the kernel line. This option will boot the system into the single user mode, which is a mode that allows the administrator to log in
as the root user and perform maintenance tasks. The GRUB menu and the kernel line are the same as the previous option. The administrator can interrupt the boot process by pressing the e key in the GRUB menu and edit the kernel line by adding systemd.unit=single.target at the end. This option will tell the system to use the single target, which is a unit that defines the state of the system in the single user mode. The administrator can then press Ctrl+X to boot the system with the modified kernel line. This option will boot the system into the single user mode and allow the administrator to troubleshoot the issues.
The other options are incorrect because they either do not boot the system into the single user mode (execute the following command from the GRUB rescue shell: mount -o remount, ro/sysroot or interrupt the boot process in the GRUB menu and add systemd.unit=single in the kernel line) or do not use the correct syntax (interrupt the boot process in the GRUB menu and add single=user in the kernel line or interrupt the boot process in the GRUB menu and add init=/bin/bash in the kernel
line). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing the Linux Boot Process, pages 267-268.

NEW QUESTION 158
A Linux system is failing to boot. The following error is displayed in the serial console: [[1;33mDEPEND[Om] Dependency failed for /data.
[[1;33mDEPEND[Om] Dependency failed for Local File Systems
...
Welcome to emergency mode! After logging in, type "journalctl -xb" to viewsystem logs,
"systemct1 reboot" to reboot, "systemct1 default" to try again to boot into default mode.
Give root password for maintenance (or type Control-D to continue}
Which of the following files will need to be modified for this server to be able to boot again?

A. /etc/mtab
B. /dev/sda
C. /etc/fstab
D. /ete/grub.conf

Answer: C

Explanation:
The file that will need to be modified for the server to be able to boot again is /etc/fstab. The /etc/fstab file is a file that contains the information about the file systems that are mounted at boot time on Linux systems. The file specifies the device name, mount point, file system type, mount options, dump frequency, and pass number for each file system. The error message indicates that the dependency failed for /data, which is a mount point for a file system. This means that the system could not mount the /data file system at boot time, which caused the system to enter the emergency mode. The emergency mode is a mode that allows the administrator to log in as the root user and perform basic tasks such as repairing the system. The administrator should modify the /etc/fstab file and check the entry for the /data file system. The administrator should look for any errors or inconsistencies in the device name, file system type, or mount options, and correct them. The administrator should also verify that the device and the file system are intact and functional by using commands such as blkid, fdisk, fsck, or mount. The administrator should then reboot the system and see if the issue is resolved. The file that will need to be modified for the server to be able to boot again is /etc/fstab. This is the correct answer to the question. The other options are incorrect because they are not related to the file systems that are mounted at boot time (/etc/mtab, /dev/sda,
or /etc/grub.conf). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 321.

NEW QUESTION 160
A Linux administrator is troubleshooting SSH connection issues from one of the workstations.
When users attempt to log in from the workstation to a server with the IP address 104.21.75.76, they receive the following message:

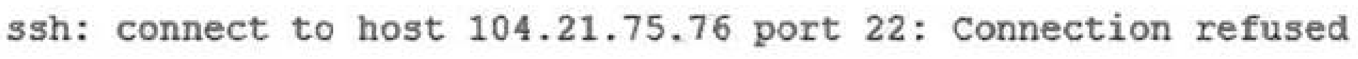

```
ssh: connect to host 104.21.75.76 port 22: Connection refused
```

The administrator reviews the information below:

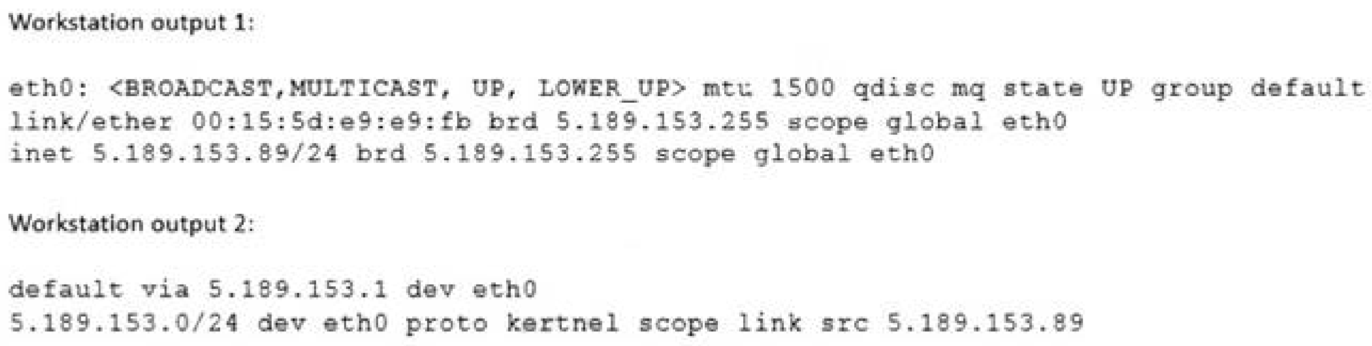

Workstation output 1:

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

Workstation output 2:

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

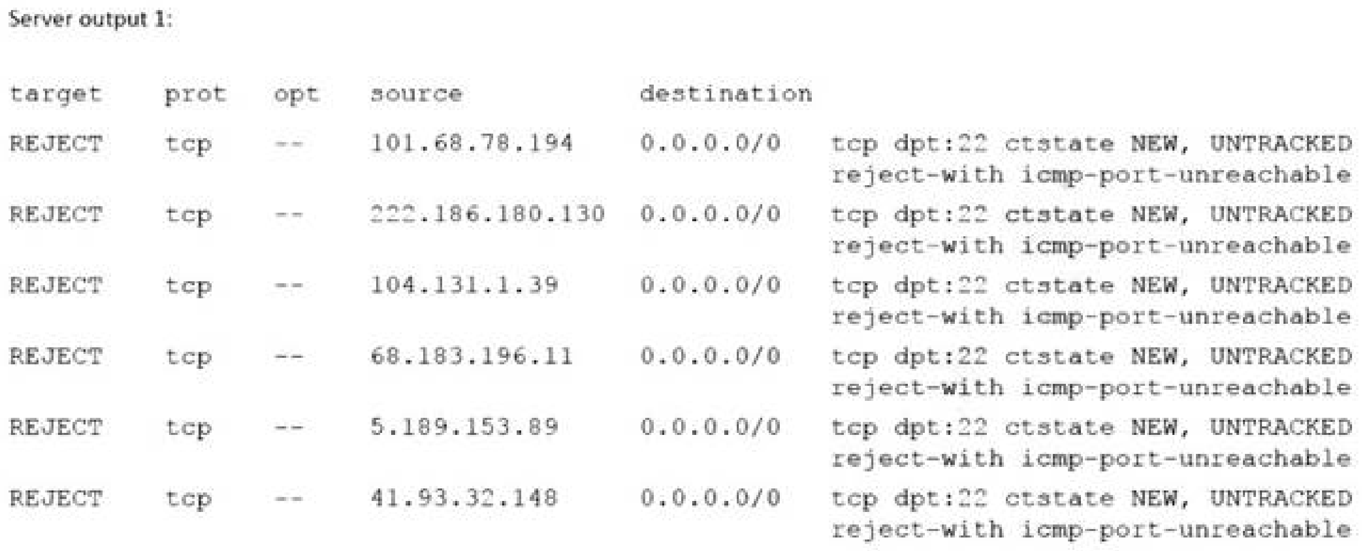

Server output 1:

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

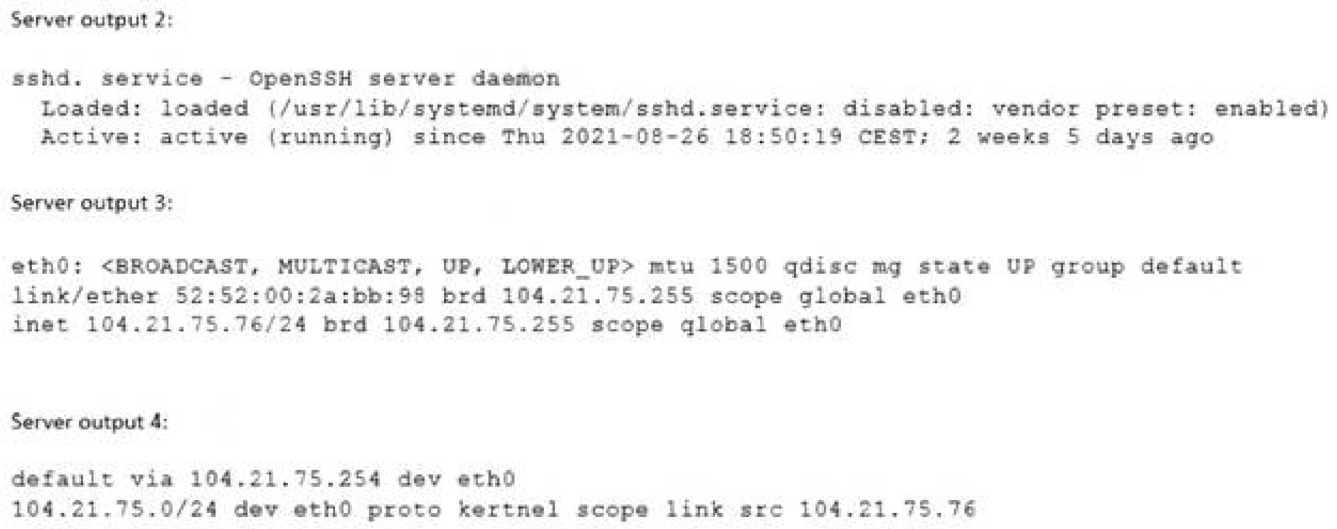

Server output 2:

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

Server output 3:

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

Server output 4:

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

Which of the following is causing the connectivity issue?

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

Answer: C

Explanation:
The server's firewall is preventing connections from being made, which is causing the connectivity issue. The output of iptables -L -n shows that the firewall is blocking all incoming traffic on port 22, which is the default port for SSH. The output of ssh -v user@104.21.75.76 shows that the connection is refused by the server. To resolve the issue, the administrator needs to allow port 22 on the firewall. The other options are incorrect because they are not supported by the outputs. The workstation has the correct IP settings, as shown by the output of ip addr show. The sshd service is enabled and running, as shown by the output of systemct1 status sshd. The server has the correct default gateway configuration, as shown by the output of ip route show. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, pages 406-407.

NEW QUESTION 165
A Linux administrator is creating a new sudo profile for the accounting user. Which of the following should be added by the administrator to the sudo configuration file so that the accounting user can run / opt/ acc/ report as root?

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

Answer: C

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

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

NEW QUESTION 166
A user is asking the systems administrator for assistance with writing a script to verify whether a file exists. Given the following:

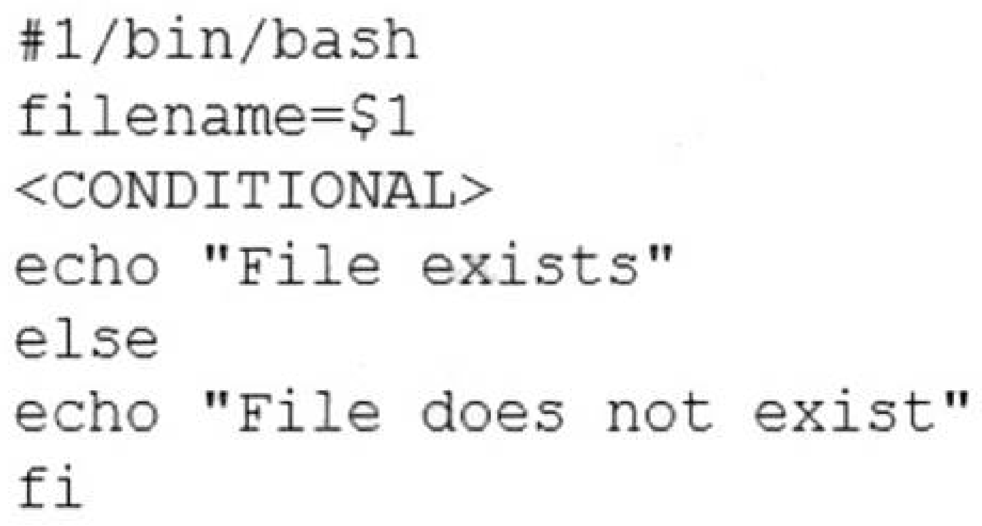

```
#1/bin/bash
filename=$1
<CONDITIONAL>
echo "File exists"
else
echo "File does not exist"
fi
```

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

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

Answer: A

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

NEW QUESTION 171
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 174
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 179
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 180
A new disk was presented to a server as /dev/ sdd. The systems administrator needs to check if a partition table is on that disk. Which of the following commands can show this information?

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

Answer: B

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

NEW QUESTION 183
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:

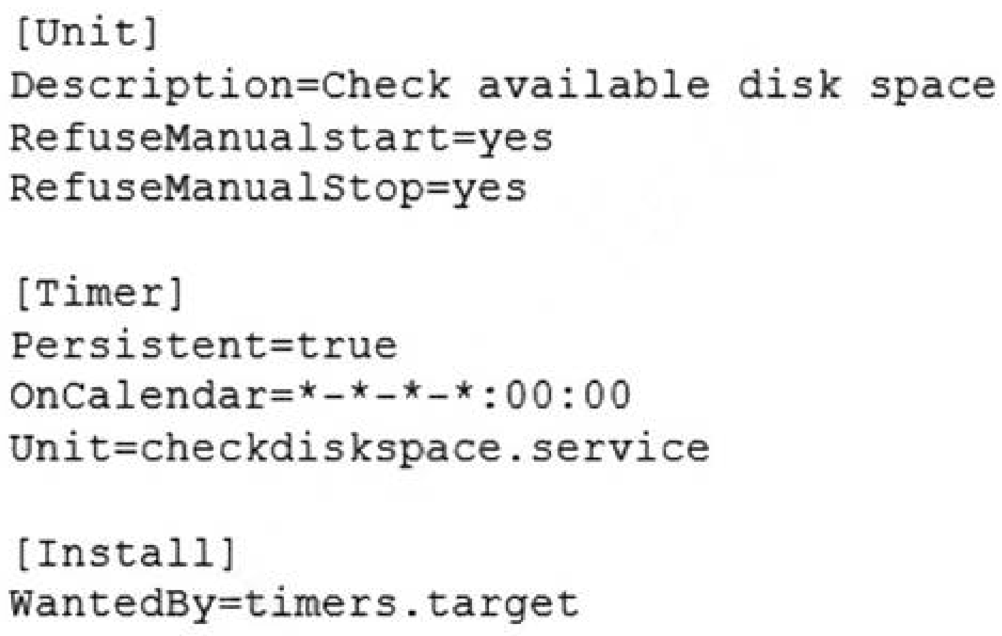

```
[Unit]
Description=Check available disk space
RefuseManualstart=yes
RefuseManualStop=yes

[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:

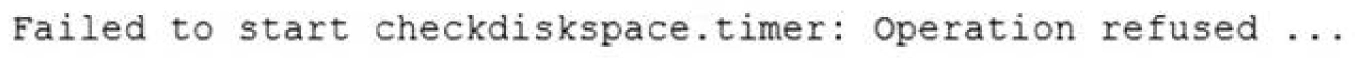

```
Failed to start checkdiskspace.timer: Operation refused ...
```

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

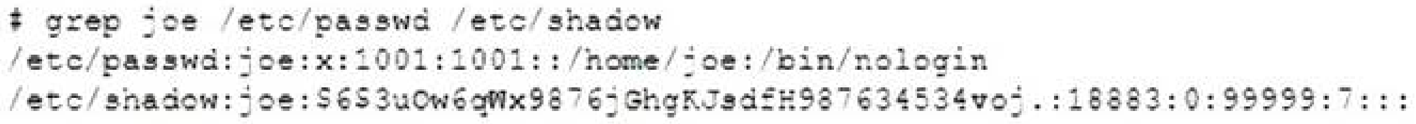

```
# grep joe /etc/passwd /etc/shadow
/etc/passwd:joe:x:1001:1001::/home/joe:/bin/nologin
/etc/shadow:joe:S6S3uOw6qWx9876jGhgKJsdfH987634534voj.:18883:0:99999:7:::
```

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 189
A systems administrator wants to test the route between IP address 10.0.2.15 and IP address 192.168.1.40. Which of the following commands will accomplish this task?

A. route -e get to 192.168.1.40 from 10.0.2.15
B. ip route get 192.163.1.40 from 10.0.2.15
C. ip route 192.169.1.40 to 10.0.2.15
D. route -n 192.168.1.40 from 10.0.2.15

Answer: B

Explanation:
The command ip route get 192.168.1.40 from 10.0.2.15 will test the route between the IP address 10.0.2.15 and the IP address 192.168.1.40. The ip route get command shows the routing decision for a given destination and source address. This is the correct command to accomplish the task. The other options are incorrect because they either use the wrong commands (route instead of ip route), the wrong options (-e or - n instead of get), or the wrong syntax (to instead of from). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 379.

NEW QUESTION 190
A cloud engineer needs to check the link status of a network interface named eth1 in a Linux server. Which of the following commands can help to achieve the goal?

A. ifconfig hw eth1
B. netstat -r eth1
C. ss -ti eth1
D. ip link show eth1

Answer: D

Explanation:
The ip link show eth1 command can be used to check the link status of a network interface named eth1 in a Linux server. It will display information such as the MAC address, MTU, state, and flags of the interface. The ifconfig hw eth1 command is invalid, as hw is not a valid option for ifconfig. The netstat -r eth1 command would display the routing table for eth1, not the link status. The ss -ti eth1 command would display TCP information for sockets associated with eth1, not the link status. References: CompTIA Linux+ (XK0- 005) Certification Study Guide, Chapter 13: Networking Fundamentals, page 436.

NEW QUESTION 192
A systems administrator wants to upgrade /bin/ someapp to a new version, but the administrator does not know the package name. Which of the following will show the RPM package name that provides that binary file?

A. rpm —qf /bin/ someapp
B. rpm —Vv / bin/ someapp
C. rpm - P / bin/ some app

D. rpm —i / bin/ someapp

Answer: A

Explanation:
The rpm command is used to manage RPM packages on Linux systems. The -qf option queries the package name that provides a given file. Therefore, the command rpm -qf /bin/someapp will show the RPM package name that provides the binary file /bin/someapp. The statements B, C, and D are incorrect because they do not query the package name, but rather verify, remove, or install a package. References: [How to Use RPM Command in Linux with Examples]

NEW QUESTION 193
A new file was added to a main Git repository. An administrator wants to synchronize a local copy with the contents of the main repository. Which of the following commands should the administrator use for this task?

A. git reflog
B. git pull
C. git status
D. git push

Answer: B

Explanation:
The command iptables -t nat -A PREROUTING -p tcp --dport 80 -j DNAT -- to-destination 192.0.2.25:3128 adds a rule to the nat table that redirects all incoming TCP packets with destination port 80 (HTTP) to the proxy server 192.0.2.25 on port 3128. This is the correct way to achieve the task. The other options are incorrect because they either delete a rule (-D), use the wrong protocol (top instead of tcp), or use the wrong port (81 instead of 80). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 381.

NEW QUESTION 198
Which of the following tools is BEST suited to orchestrate a large number of containers across many different servers?

A. Kubernetes
B. Ansible
C. Podman
D. Terraform

Answer: A

Explanation:
The tool that is best suited to orchestrate a large number of containers across many different servers is Kubernetes. Kubernetes is an open-source platform for managing containerized applications and services. Kubernetes allows the administrator to deploy, scale, and update containers across a cluster of servers, as well as to automate the configuration and coordination of the containers. Kubernetes also provides features such as service discovery, load balancing, storage management, security, monitoring, and logging. Kubernetes can handle complex and dynamic workloads and ensure high availability and performance of the containers. Kubernetes is the tool that is best suited to orchestrate a large number of containers across many different servers. This is the correct answer to the question. The other options are incorrect because they either do not orchestrate containers (Ansible or Terraform) or do not operate across many different servers (Podman). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 573.

NEW QUESTION 202
A Linux administrator was tasked with deleting all files and directories with names that are contained in the sobelete.txt file. Which of the following commands will accomplish this task?

A. xargs -f cat toDelete.txt -rm
B. rm -d -r -f toDelete.txt
C. cat toDelete.txt | rm -frd
D. cat toDelete.txt | xargs rm -rf

Answer: D

Explanation:
The command cat toDelete.txt | xargs rm -rf will delete all files and directories with names that are contained in the toDelete.txt file. The cat command reads the file and outputs its contents to the standard output. The | operator pipes the output to the next command. The xargs command converts the output into arguments for the next command. The rm -rf command removes the files and directories recursively and forcefully. This is the correct way to accomplish the task. The other options are incorrect because they either use the wrong options (-f instead of -a for xargs), the wrong arguments (toDelete.txt instead of toDelete.txt filename for rm), or the wrong commands (rm instead of xargs). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, pages 349-350.

NEW QUESTION 203
A server is experiencing intermittent connection issues. Some connections to the Internet work as intended, but some fail as if there is no connectivity. The systems administrator inspects the server configuration:

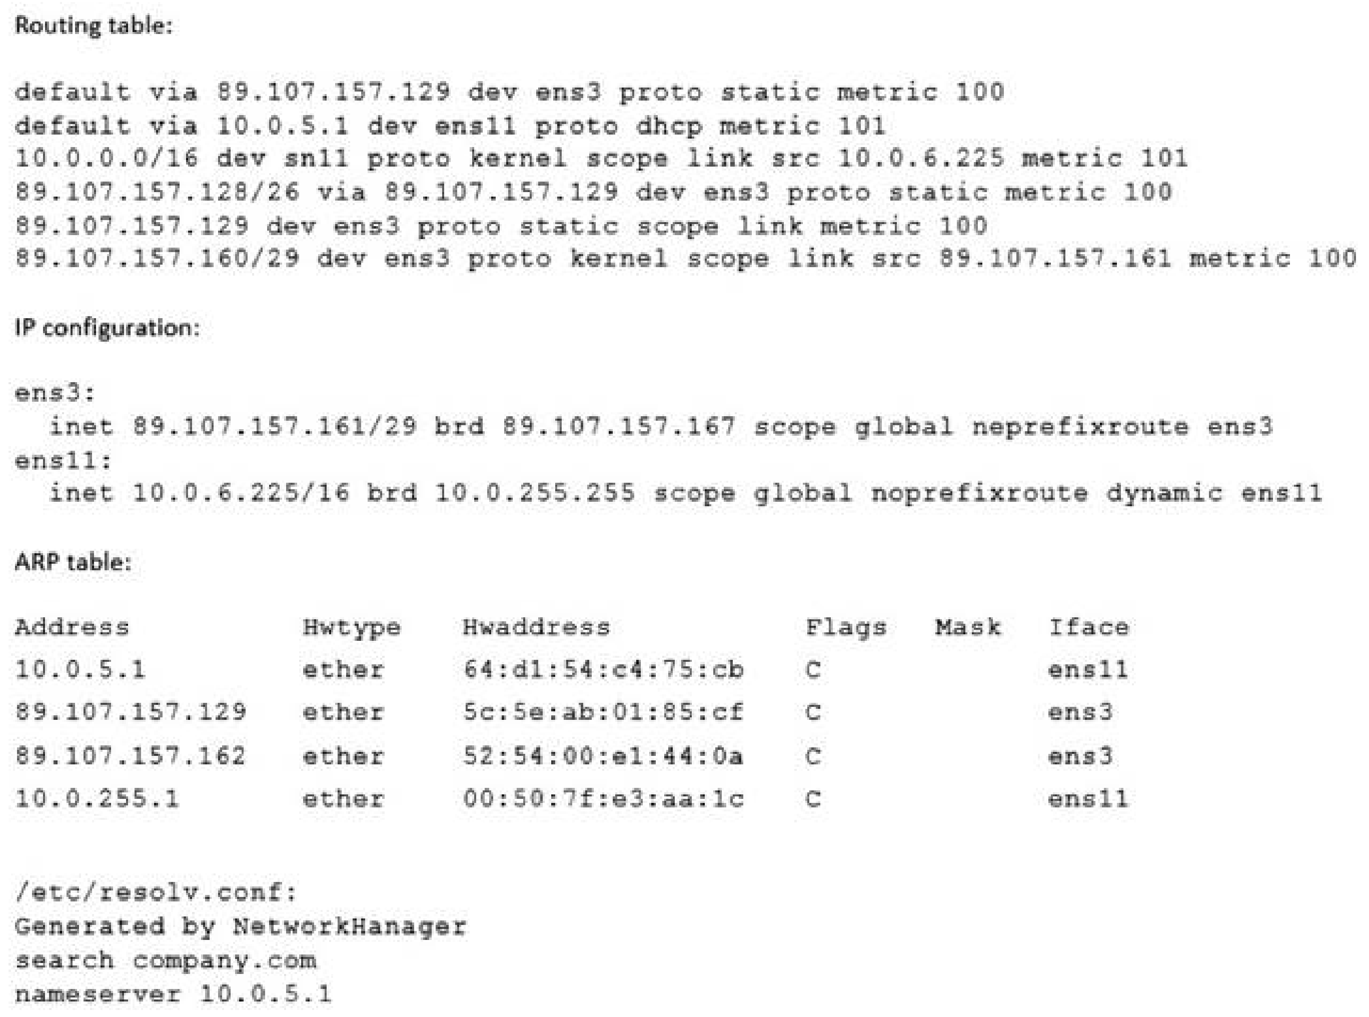

Routing table:

```
default via 89.107.157.129 dev ens3 proto static metric 100
default via 10.0.5.1 dev ens11 proto dhcp metric 101
10.0.0.0/16 dev sn11 proto kernel scope link src 10.0.6.225 metric 101
89.107.157.128/26 via 89.107.157.129 dev ens3 proto static metric 100
89.107.157.129 dev ens3 proto static scope link metric 100
89.107.157.160/29 dev ens3 proto kernel scope link src 89.107.157.161 metric 100
```

IP configuration:

```
ens3:
  inet 89.107.157.161/29 brd 89.107.157.167 scope global neprefixroute ens3
ens11:
  inet 10.0.6.225/16 brd 10.0.255.255 scope global noprefixroute dynamic ens11
```

ARP table:

```
Address           Hwtype    Hwaddress            Flags   Mask   Iface
10.0.5.1          ether     64:d1:54:c4:75:cb    C              ens11
89.107.157.129    ether     5c:5e:ab:01:85:cf    C              ens3
89.107.157.162    ether     52:54:00:e1:44:0a    C              ens3
10.0.255.1        ether     00:50:7f:e3:aa:1c    C              ens11
```

```
/etc/resolv.conf:
Generated by NetworkHanager
search company.com
nameserver 10.0.5.1
```

Which of the following is MOST likely the cause of the issue?

A. An internal-only DNS server is configured.
B. The IP netmask is wrong for ens3.
C. Two default routes are configured.
D. The ARP table contains incorrect entries.

Answer: C

Explanation:
The most likely cause of the issue is that two default routes are configured on the server. The default route is the route that is used when no other route matches the destination of a packet. The default route is usually the gateway that connects the local network to the Internet. The server configuration shows that there are two default routes in the routing table, one with the gateway 192.168.1.1 and the other with the gateway 10.0.0.1. This can cause a conflict and confusion for the server when deciding which gateway to use for the outgoing packets. Some packets may be sent to the wrong gateway and fail to reach the Internet, while some packets may be sent to the correct gateway and work as intended. This can result in intermittent connection issues and inconsistent behavior. The administrator should remove one of the default routes and keep only the correct one for the network. This can be done by using the ip route del command or by editing the network configuration files. This will resolve the issue and restore the connectivity. The other options are incorrect because they are not supported by the outputs. The DNS server, the IP netmask, and the ARP table are not the causes of the issue. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, pages 381-382.

NEW QUESTION 204
Which of the following should be used to verify the integrity of a file?

A. sha256sum
B. fsck
C. gpg —d
D. hashcat

Answer: A

Explanation:
The best tool to use to verify the integrity of a file is A. sha256sum. This tool will compute and display the SHA-256 hash of a file, which is a 64-digit hexadecimal number that uniquely identifies the file's content. By comparing the hash of a downloaded file with the hash provided by the file owner or source, you can confirm that the file has not been altered or corrupted during the transfer. The other tools are either not relevant or not suitable for this task. For example:
? B. fsck is a tool for checking and repairing the file system, but it does not verify the
integrity of individual files.
? C. gpg -d is a tool for decrypting files that have been encrypted with GnuPG, but it does not verify the integrity of unencrypted files.
? D. hashcat is a tool for cracking passwords or hashes, but it does not verify the integrity of files.

NEW QUESTION 207
A junior administrator updated the PostgreSQL service unit file per the data-base administrator's recommendation. The service has been restarted, but changes have not been applied. Which of the following should the administrator run for the changes to take effect?

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

Answer: B

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

NEW QUESTION 211
Which of the following commands will display the operating system?

A. uname -n
B. uname -s
C. uname -o
D. uname -m

Answer: C

Explanation:
The command that will display the operating system is uname -o. This command uses the uname tool, which is used to print system information such as the kernel name, version, release, machine, and processor. The -o option stands for operating system, and prints the name of the operating system implementation (usually GNU/Linux). The other options are not correct commands for displaying the operating system. The uname -n command will display the network node hostname of the system. The uname -s command will display the kernel name of the system. The uname -m command will display the machine hardware name of the system. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 1: Exploring Linux Command-Line Tools; uname(1) - Linux manual page

NEW QUESTION 215
A systems administrator pressed Ctrl+Z after starting a program using the command line, and the shell prompt was presented. In order to go back to the program, which of the following commands can the administrator use?

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

Answer: A

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

NEW QUESTION 219
A Linux administrator created the directory /project/access2all. By creating this directory, the administrator is trying to avoid the deletion or modification of files from non-owners. Which of the following will accomplish this goal?

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

Answer: A

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

NEW QUESTION 223
DRAG DROP
As a Systems Administrator, to reduce disk space, you were tasked to create a shell script that does the following:
Add relevant content to /tmp/script.sh, so that it finds and compresses rotated files in /var/log without recursion. INSTRUCTIONS
Fill the blanks to build a script that performs the actual compression of rotated log files.
If at any time you would like to bring back the initial state of the simulation, please click the Reset All button.

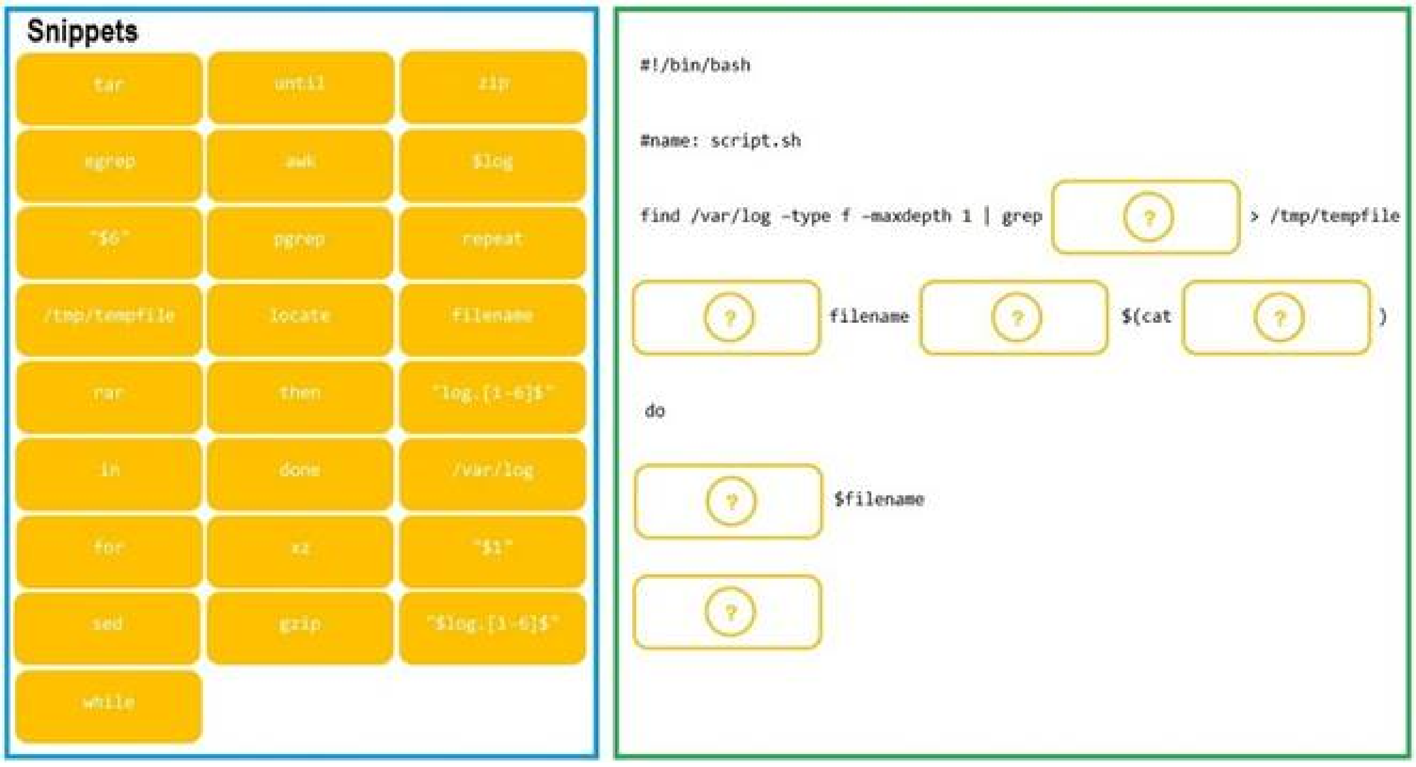

```
#!/bin/bash

#name: script.sh

find /var/log -type f -maxdepth 1 | grep [?] > /tmp/tempfile

[?] filename [?] $(cat [?] )

do

[?] $filename

[?]
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:

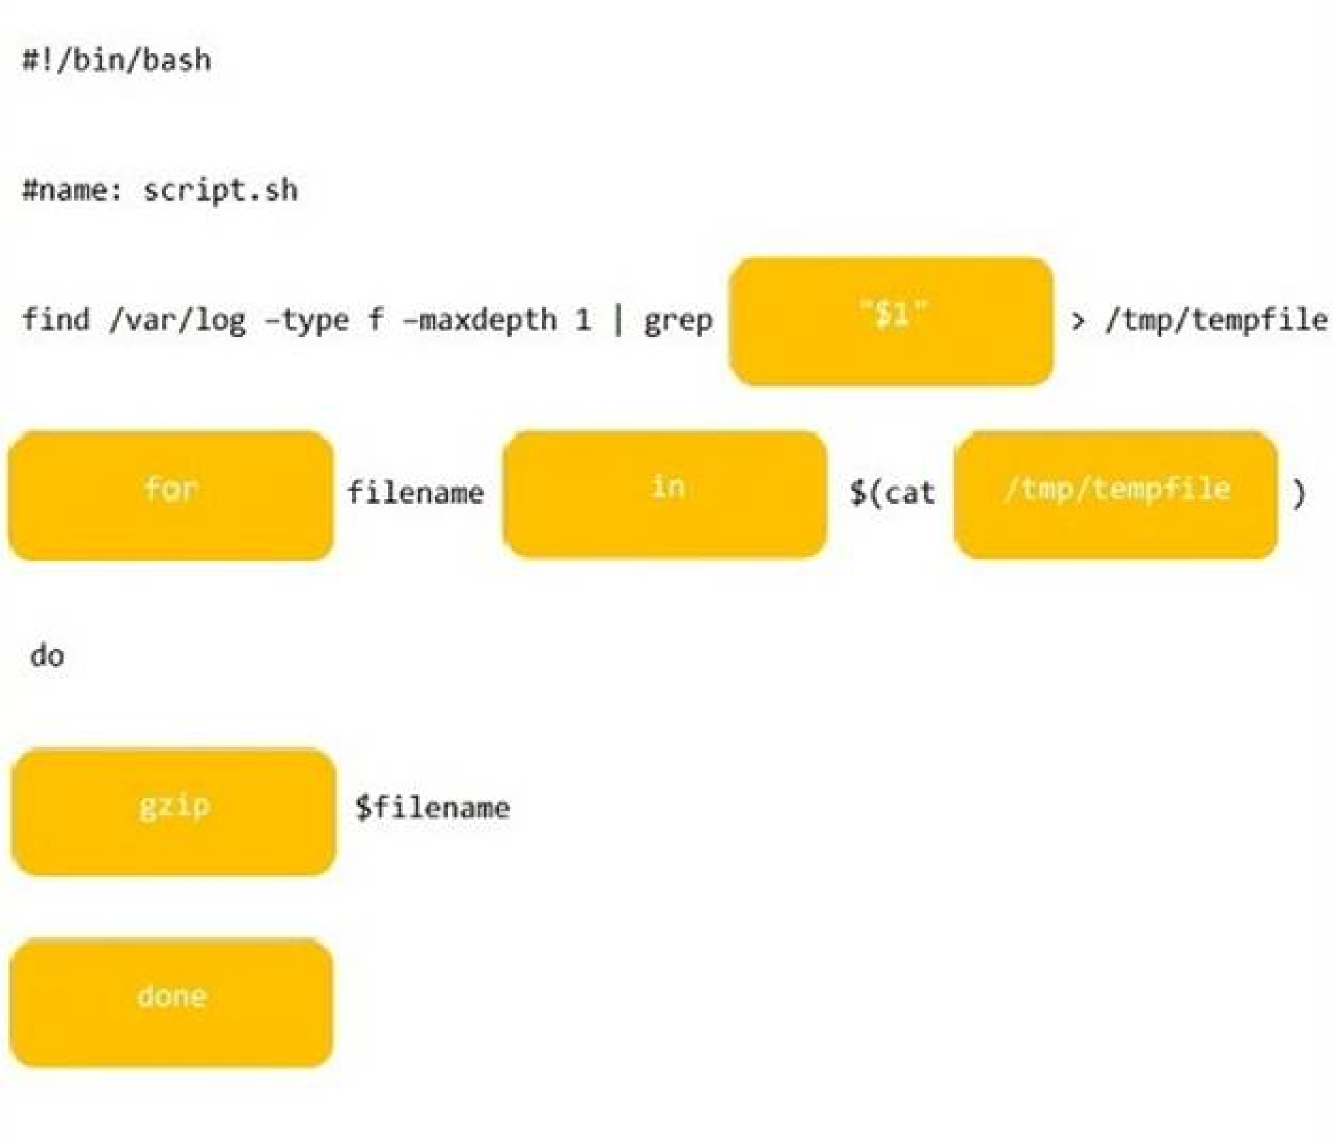

```
#!/bin/bash

#name: script.sh

find /var/log -type f -maxdepth 1 | grep "$1" > /tmp/tempfile

for filename in $(cat /tmp/tempfile )

do

gzip $filename

done
```

NEW QUESTION 227
......