CompTIA

Exam Questions XK0-005

CompTIA Linux+ Certification Exam

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

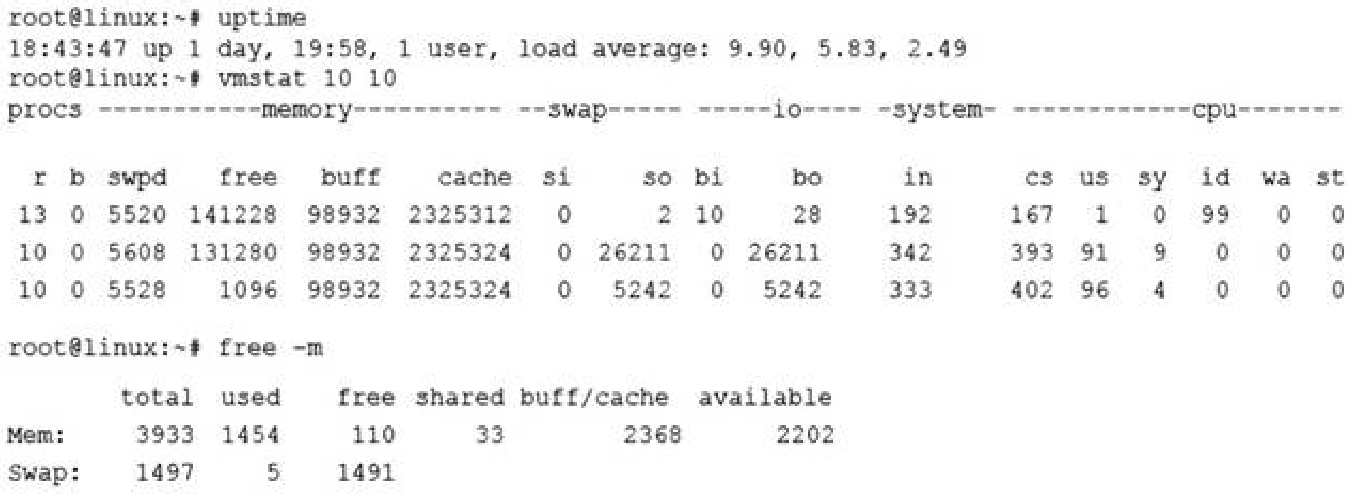

```
root@linux:~# uptime
18:43:47 up 1 day, 19:58, 1 user, load average: 9.90, 5.83, 2.49
root@linux:~# vmstat 10 10
procs -----------memory---------- --swap----- -----io---- -system- ------------cpu-------

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

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

Answer: D

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

NEW QUESTION 4
A Linux administrator is configuring a new internal web server fleet. The web servers are up and running but can only be reached by users directly via IP address. The administrator is attempting to fix this inconvenience by requesting appropriate records from the DNS team. The details are:
Hostname: devel.comptia.org
IP address: 5.5.5.1, 5.5.5.2, 5.5.5.3, 5.5.5.4
Name server: 5.5.5.254

Additional names: dev.comptia.org, development.comptia.org
Which of the following types of DNS records should the Linux administrator request from the DNS team? (Select three).

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

Answer: BDE

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

NEW QUESTION 5
An administrator has source code and needs to rebuild a kernel module. Which of the following command sequences is most commonly used to rebuild this type of module?

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

Answer: A

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

NEW QUESTION 6
Application code is stored in Git. Due to security concerns, the DevOps engineer does not want to keep a sensitive configuration file, app . conf, in the repository. Which of the following should the engineer do to prevent the file from being uploaded to the repository?

A. Run git exclude ap
B. conf.
C. Run git stash ap
D. conf.
E. Add app . conf to . exclude.
F. Add app . conf to . gitignore.

Answer: D

Explanation:
This will prevent the file app.conf from being tracked by Git and uploaded to the repository. The .gitignore file is a special file that contains patterns of files and directories that Git should ignore. Any file that matches a pattern in the .gitignore file will not be staged, committed, or pushed to the remote repository. The .gitignore file should be placed in the root directory of the repository and committed along with the other files.
The other options are incorrect because:
* A. Run git exclude app.conf
This is not a valid Git command. There is no such thing as git exclude. The closest thing is git update-index --assume-unchanged, which tells Git to temporarily ignore changes to a file, but it does not prevent the file from being uploaded to the repository.
* B. Run git stash app.conf
This will temporarily save the changes to the file app.conf in a stash, which is a hidden storage area for uncommitted changes. However, this does not prevent the

file from being tracked by Git or uploaded to the repository. The file will still be part of the working tree and the index, and it will be restored when the stash is popped or applied.
* C. Add app.conf to .exclude
This will have no effect, because Git does not recognize a file named .exclude. The only files that Git uses to ignore files are .gitignore, $GIT_DIR/info/exclude, and core.excludesFile.
References:
? Git - gitignore Documentation
? .gitignore file - ignoring files in Git | Atlassian Git Tutorial
? Ignoring files - GitHub Docs
? [CompTIA Linux+ Certification Exam Objectives]

NEW QUESTION 7
A cloud engineer needs to block the IP address 192.168.10.50 from accessing a Linux server. Which of the following commands will achieve this goal?

A. iptables -F INPUT -j 192.168.10.50 -m DROP
B. iptables -A INPUT -s 192.168.10.30 -j DROP
C. iptables -i INPUT --ipv4 192.168.10.50 -z DROP
D. iptables -j INPUT 192.168.10.50 -p DROP

Answer: B

Explanation:
The correct command to block the IP address 192.168.10.50 from accessing a Linux server is iptables -A INPUT -s 192.168.10.50 -j DROP. This command appends a rule to the INPUT chain that matches the source address 192.168.10.50 and jumps to the DROP target, which discards the packet. The other commands are incorrect because they either have invalid syntax, wrong parameters, or wrong order of arguments. References:
CompTIA Linux+ Study Guide, Fourth Edition, page 457-458.

NEW QUESTION 8
A systems administrator is tasked with setting up key-based SSH authentication. In which of the following locations should the administrator place the public keys for the server?

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

Answer: C

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

NEW QUESTION 9
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 10
A systems administrator frequently connects to a remote host via SSH and a non-standard port. The systems administrator would like to avoid passing the port parameter on the command line every time. Which of the following files can be used to set a different port
value for that host?

A. /etc/ssh/sshd_config
B. /etc/ssh/moduli
C. ~/.ssh/config
D. ~/.ssh/authorized_keys

Answer: C

Explanation:
The ~/.ssh/config file can be used to set various options for SSH connections, including the port number, for specific hosts or groups of hosts. This file is located in the user's home directory and affects only the current user. The /etc/ssh/sshd_config file is used to configure the SSH server daemon, not the client. The /etc/ssh/moduli file contains parameters for Diffie-Hellman key exchange, not port settings.
The ~/.ssh/authorized_keys file contains public keys for authentication, not port settings. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Secure Shell (SSH), page 414.

NEW QUESTION 10
A systems administrator needs to clone the partition /dev/sdc1 to /dev/sdd1. Which of the following commands will accomplish this task?

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

Answer: C

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

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

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

Answer: C

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

NEW QUESTION 21
A systems administrator requires that all files that are created by the user named web have read-only permissions by the owner. Which of the following commands will satisfy this requirement?

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

Answer: C

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

NEW QUESTION 26
Which of the following directories is the mount point in a UEFI system?

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

Answer: B

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

NEW QUESTION 30
A User on a Linux workstation needs to remotely start an application on a Linux server and then forward the graphical display of that application back to the Linux workstation. Which of the following would enable the user to perform this action?

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

Answer: A

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

NEW QUESTION 32
A Linux engineer needs to block an incoming connection from the IP address 2.2.2.2 to a secure shell server and ensure the originating IP address receives a response that a firewall is blocking the connection. Which of the following commands can be used to accomplish this task?

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

Answer: C

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

NEW QUESTION 36
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 39
A Linux system fails to start and delivers the following error message:

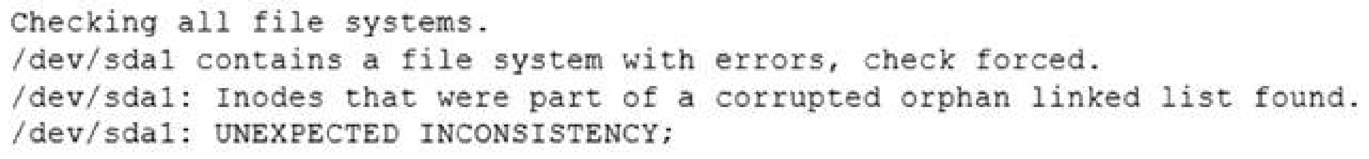

```
Checking all file systems.
/dev/sda1 contains a file system with errors, check forced.
/dev/sda1: Inodes that were part of a corrupted orphan linked list found.
/dev/sda1: UNEXPECTED INCONSISTENCY;
```

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 43
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 47
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 48
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 52
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 56
A systems administrator is enabling LUKS on a USB storage device with an ext4 filesystem format. The administrator runs dmesg and notices the following output:

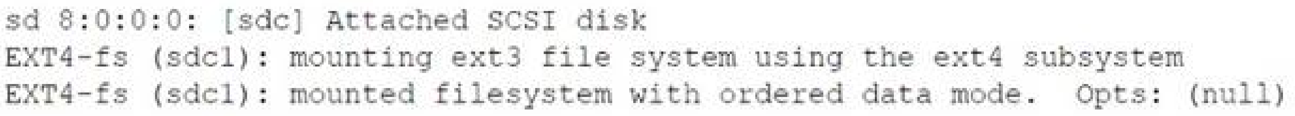

```
sd 8:0:0:0: [sdc] Attached SCSI disk
EXT4-fs (sdc1): mounting ext3 file system using the ext4 subsystem
EXT4-fs (sdc1): mounted filesystem with ordered data mode.  Opts: (null)
```

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

A. gpg /dev/sdcl
B. pvcreate /dev/sdc
C. mkfs . ext4 /dev/mapper/LUKSCJ001 - L ENCRYPTED
D. umount / dev/ sdc
E. fdisk /dev/sdc
F. mkfs . vfat /dev/mapper/LUKS0001 — L ENCRYPTED
G. wipefs —a/dev/sdbl
H. cryptsetup luksFormat /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 61
Which of the following tools is commonly used for creating CI/CD pipelines?

A. Chef
B. Puppet
C. Jenkins
D. Ansible

Answer: C

Explanation:
The tool that is commonly used for creating CI/CD pipelines is Jenkins. Jenkins is an open-source automation server that enables continuous integration and continuous delivery (CI/CD) of software projects. Jenkins allows developers to build, test, and deploy code changes automatically and frequently using various plugins and integrations. Jenkins also supports distributed builds, parallel execution, pipelines as code, and real-time feedback. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Source Code; Jenkins

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

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

Answer: A

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

NEW QUESTION 67
A junior administrator is setting up a new Linux server that is intended to be used as a router at a remote site. Which of the following parameters will accomplish this goal?
A.

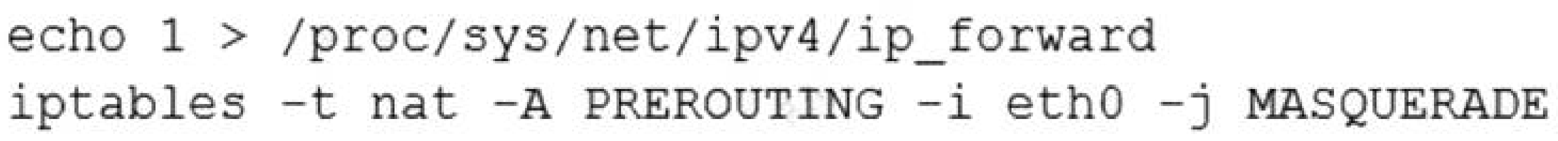

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A PREROUTING -i eth0 -j MASQUERADE
```

A.

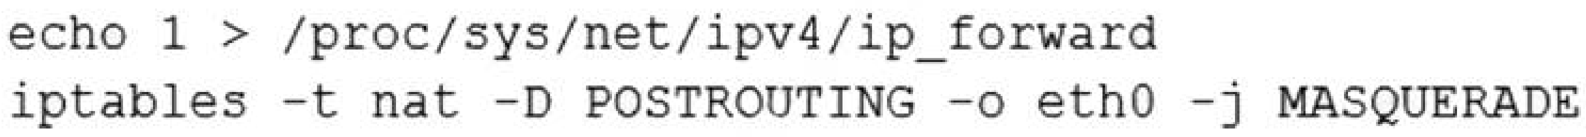

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -D POSTROUTING -o eth0 -j MASQUERADE
```

B.

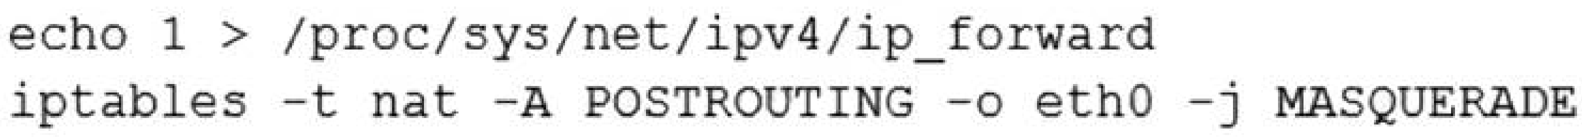

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
```

C.

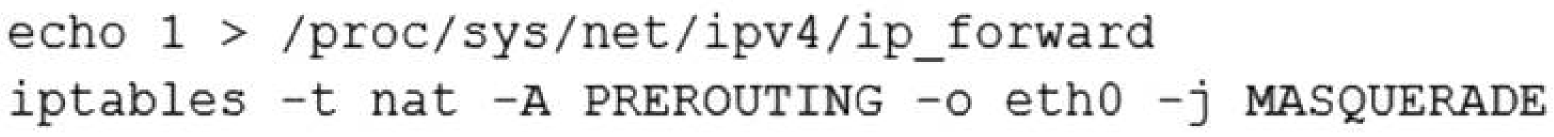

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A PREROUTING -o eth0 -j MASQUERADE
```

Answer: C

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

NEW QUESTION 71
Due to performance issues on a server, a Linux administrator needs to termi-nate an unresponsive process. Which of the following commands should the administrator use to terminate the process immediately without waiting for a graceful shutdown?

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

Answer: A

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

NEW QUESTION 74
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 78
A Linux engineer has been notified about the possible deletion of logs from the file
/opt/app/logs. The engineer needs to ensure the log file can only be written into without removing previous entries.

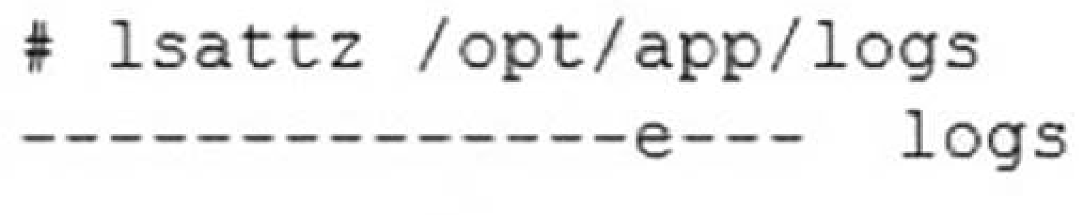

```
# lsattz /opt/app/logs
--------------e---   logs
```

Which of the following commands would be BEST to use to accomplish this task?

A. chattr +a /opt/app/logs
B. chattr +d /opt/app/logs
C. chattr +i /opt/app/logs
D. chattr +c /opt/app/logs

Answer: A

Explanation:
The command chattr +a /opt/app/logs will ensure the log file can only be written into without removing previous entries. The chattr command is a tool for changing file attributes on Linux file systems. The +a option sets the append-only attribute, which means that the file can only be opened in append mode for writing. This prevents the file from being modified, deleted, or renamed. This is the best command to use to accomplish the task. The other options are incorrect because they either set the wrong attributes
(+d, +i, or +c) or do not affect the file at all (-a). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, page 357.

NEW QUESTION 80
A Linux administrator needs to determine whether a hostname is in the DNS. Which of the following would supply the information that is needed?

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

Answer: A

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

NEW QUESTION 83
Users are unable to create new files on the company's FTP server, and an administrator is troubleshooting the issue. The administrator runs the following commands:

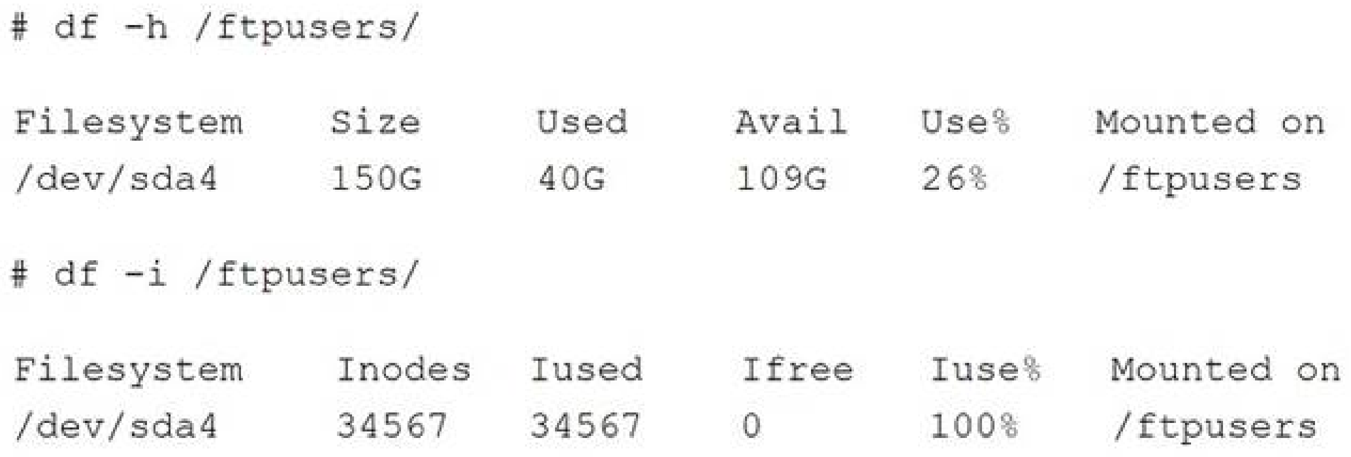

```
# df -h /ftpusers/

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

# df -i /ftpusers/

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

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

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

Answer: C

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

NEW QUESTION 86
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 91
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 94
A Linux administrator is creating a primary partition on the replacement hard drive for an application server. Which of the following commands should the administrator issue to verify the device name of this partition?

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

Answer: C

Explanation:
The command sudo fdisk -l should be issued to verify the device name of the partition. The sudo command allows the administrator to run commands as the superuser or another user. The fdisk command is a tool for manipulating disk partitions on Linux systems. The -l option lists the partitions on all disks or a specific disk. The command sudo fdisk -l will show the device names, sizes, types, and other information of the partitions on all disks. The administrator can identify the device name of the partition by looking at the output. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not list the partitions (sudo fdisk /dev/sda or sudo fdisk -h) or do not exist (sudo fdisk -s /dev/sda). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 317.

NEW QUESTION 97
A systems administrator received a request to change a user's credentials. Which of the following commands will grant the request?

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

Answer: A

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

NEW QUESTION 98
A Linux administrator found many containers in an exited state. Which of the following commands will allow the administrator to clean up the containers in an exited state?

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

Answer: B

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

NEW QUESTION 100
One leg of an LVM-mirrored volume failed due to the underlying physical volume, and a systems administrator is troubleshooting the issue. The following output has been provided:

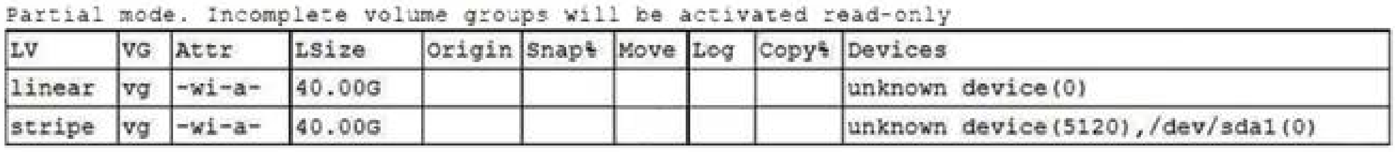

```
Partial mode. Incomplete volume groups will be activated read-only
```

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

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

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

Answer: B

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

NEW QUESTION 104
A Linux engineer needs to download a ZIP file and wants to set the nice of value to -10 for this new process. Which of the following commands will help to accomplish the task?

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

Answer: D

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

NEW QUESTION 108
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 112
A systems administrator needs to remove a disk from a Linux server. The disk size is 500G, and it is the only one that size on that machine. Which of the following commands can the administrator use to find the corresponding device name?

A. fdisk -V
B. partprobe -a
C. lsusb -t
D. lsscsi -s

Answer: D

Explanation:
The lsscsi command can list the SCSI devices on the system, along with their size and device name. The -s option shows the size of each device. The administrator can look for the device that has a size of 500G and note its device name. See lsscsi(8) - Linux man page and How to check Disk Interface Types in Linux.References1: https://linux.die.net/man/8/lsscsi2: https://www.golinuxcloud.com/check-disk-type-linux/

NEW QUESTION 113
A systems administrator is trying to track down a rogue process that has a TCP listener on a network interface for remote command-and-control instructions. Which of the following commands should the systems administrator use to generate a list of rogue process names? (Select two).

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

Answer: AB

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

NEW QUESTION 116
An administrator transferred a key for SSH authentication to a home directory on a remote server. The key file was moved to .ssh/authorized_keys location in order to establish SSH connection without a password. However, the SSH command still asked for the password. Given the following output:

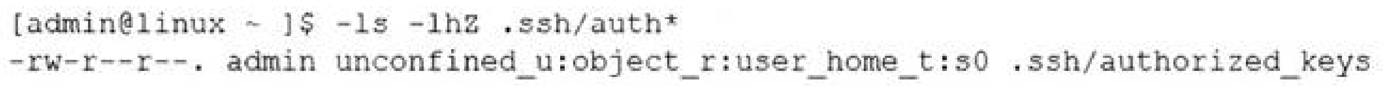

```
[admin@linux ~ ]$ -ls -lhZ .ssh/auth*
-rw-r--r--. admin unconfined_u:object_r:user_home_t:s0 .ssh/authorized_keys
```

Which of the following commands would resolve the issue?

A. restorecon .ssh/authorized_keys
B. ssh_keygen -t rsa -o .ssh/authorized_keys
C. chown root:root .ssh/authorized_keys
D. chmod 600 .ssh/authorized_keys

Answer: D

Explanation:
The command that would resolve the issue is chmod 600 .ssh/authorized_keys. This command will change the permissions of the .ssh/authorized_keys file to 600, which means that only the owner of the file can read and write it. This is necessary for SSH key authentication to work properly, as SSH will refuse to use a key file that is accessible by other users or groups for security reasons. The output of ls -l shows that currently the .ssh/authorized_keys file has permissions of 664, which means that both the owner and group can read and write it, and others can read it.
The other options are not correct commands for resolving the issue. The restorecon .ssh/authorized_keys command will restore the default SELinux security context for the .ssh/authorized_keys file, but this will not change its permissions or ownership. The ssh_keygen -t rsa -o .ssh/authorized_keys command is invalid because ssh_keygen is not a valid command (the correct command is ssh-keygen), and the -o option is used to specify a new output format for the key file, not the output file name. The chown root:root
.ssh/authorized_keys command will change the owner and group of the .ssh/authorized_keys file to root, but this will not change its permissions or make it accessible by the user who wants to log in with SSH key authentication. References: How to Use Public Key Authentication with SSH; chmod(1) - Linux manual page

NEW QUESTION 121
A Linux administrator is troubleshooting an issue in which users are not able to access https://portal.comptia.org from a specific workstation. The administrator runs a few commands and receives the following output:

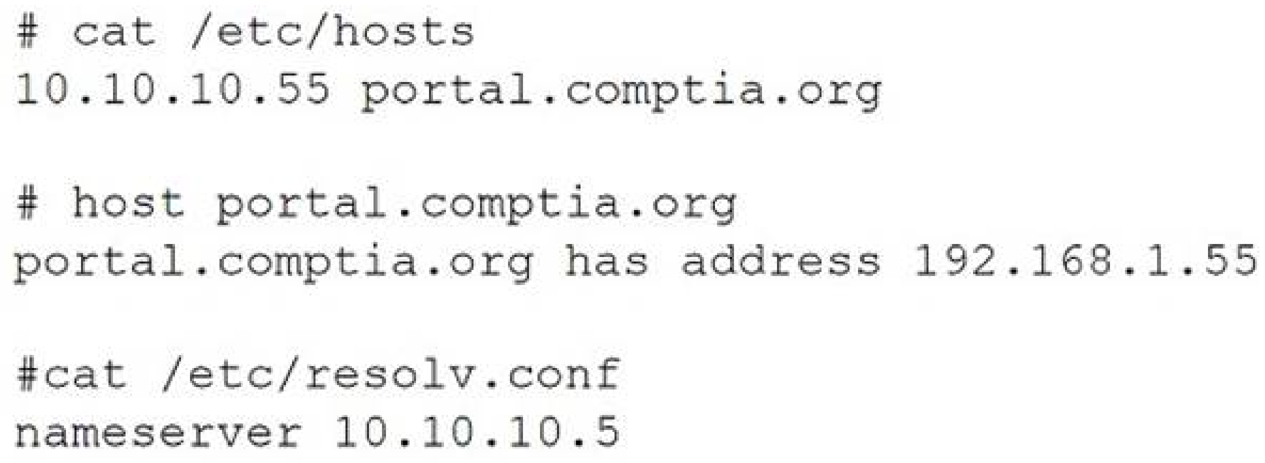

```
# cat /etc/hosts
10.10.10.55 portal.comptia.org

# host portal.comptia.org
portal.comptia.org has address 192.168.1.55

#cat /etc/resolv.conf
nameserver 10.10.10.5
```

Which of the following tasks should the administrator perform to resolve this issue?

A. Update the name server in resol
B. conf to use an external DNS server.
C. Remove the entry for portal . comptia.org from the local hosts file.
D. Add a network route from the 10.10.10.0/24 to the 192.168.0.0/16.
E. Clear the local DNS cache on the workstation and rerun the host command.

Answer: B

Explanation:
The best task to perform to resolve this issue is B. Remove the entry for portal.comptia.org from the local hosts file. This is because the local hosts file has a wrong entry that maps portal.comptia.org to 10.10.10.55, which is different from the actual IP address of 192.168.1.55 that is returned by the DNS server. This causes a mismatch and prevents the workstation from accessing the website. By removing or correcting the entry in the hosts file, the workstation will use the DNS server to resolve the domain name and access the website successfully.
To remove or edit the entry in the hosts file, you need to have root privileges and use a text editor such as vi or nano. For example, you can run the command:
sudo vi /etc/hosts
and delete or modify the line that says: 10.10.10.55 portal.comptia.org
Then save and exit the file.

NEW QUESTION 124
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 125
A Linux administrator needs to create a new cloud.cpio archive containing all the files from the current directory. Which of the following commands can help to accomplish this task?

A. ls | cpio -iv > cloud.epio
B. ls | cpio -iv < cloud.epio
C. ls | cpio -ov > cloud.cpio
D. ls cpio -ov < cloud.cpio

Answer: C

Explanation:
The command ls | cpio -ov > cloud.cpio can help to create a new cloud.cpio archive containing all the files from the current directory. The ls command lists the files in the current directory and outputs them to the standard output. The | operator pipes the output to the next command. The cpio command is a tool for creating and extracting compressed archives. The -o option creates a new archive and the -v option shows the verbose output. The > operator redirects the output to the cloud.cpio file. This command will create a new cloud.cpio archive with all the files from the current directory. The other options are incorrect because they either use the wrong options (-i instead of -o), the wrong arguments (cloud.epio instead of cloud.cpio), or the wrong syntax (< instead of > or missing |). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, page 351.

NEW QUESTION 128
A Linux administrator needs to expand a volume group using a new disk. Which of the following options presents the correct sequence of commands to accomplish the task?

A. partprobe vgcreate lvextend
B. lvcreate fdisk partprobe
C. fdisk partprobe mkfs
D. fdisk pvcreate vgextend

Answer: D

Explanation:
The correct sequence of commands to expand a volume group using a new disk is fdisk, pvcreate, vgextend. The fdisk command can be used to create a partition on the new disk with the type 8e (Linux LVM). The pvcreate command can be used to initialize the partition as a physical volume for LVM. The vgextend command can be used to add the physical volume to an existing volume group. The partprobe command can be used to inform the kernel about partition table changes, but it is not necessary in this case. The vgcreate command can be used to create a new volume group, not expand an existing one. The lvextend command can be used to extend a logical volume, not a volume group. The lvcreate command can be used to create a new logical volume, not expand a volume group. The mkfs command can be used to create a filesystem on a partition or a logical volume, not expand a volume group. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Disk Storage, pages 462-463.

NEW QUESTION 133
A systems administrator is tasked with creating a cloud-based server with a public IP address.

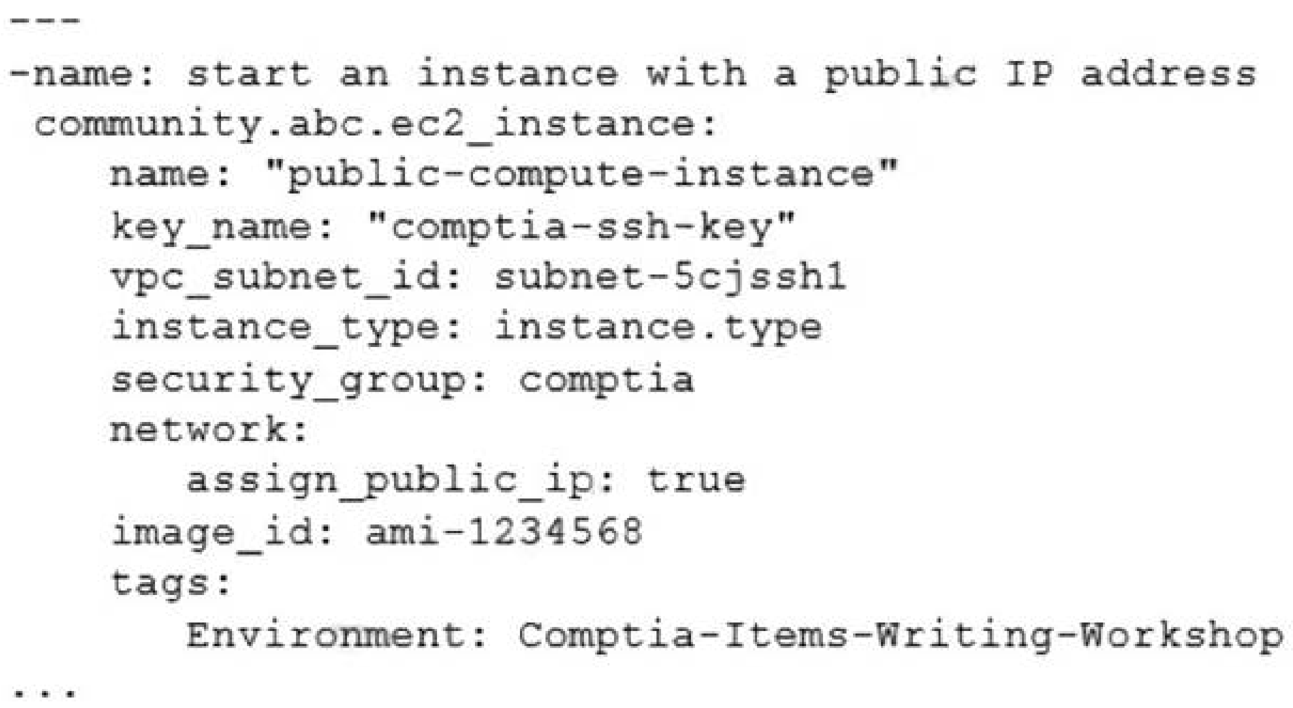

```
---
-name: start an instance with a public IP address
 community.abc.ec2_instance:
    name: "public-compute-instance"
    key_name: "comptia-ssh-key"
    vpc_subnet_id: subnet-5cjssh1
    instance_type: instance.type
    security_group: comptia
    network:
       assign_public_ip: true
    image_id: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

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

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

Answer: D

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

NEW QUESTION 137
An administrator started a long-running process in the foreground that needs to continue without interruption. Which of the following keystrokes should the administrator use to continue running the process in the background?

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

Answer: A

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

NEW QUESTION 138
An administrator added the port 2222 for the SSH server on myhost and restarted the SSH server. The administrator noticed issues during the startup of the

service. Given the following outputs:

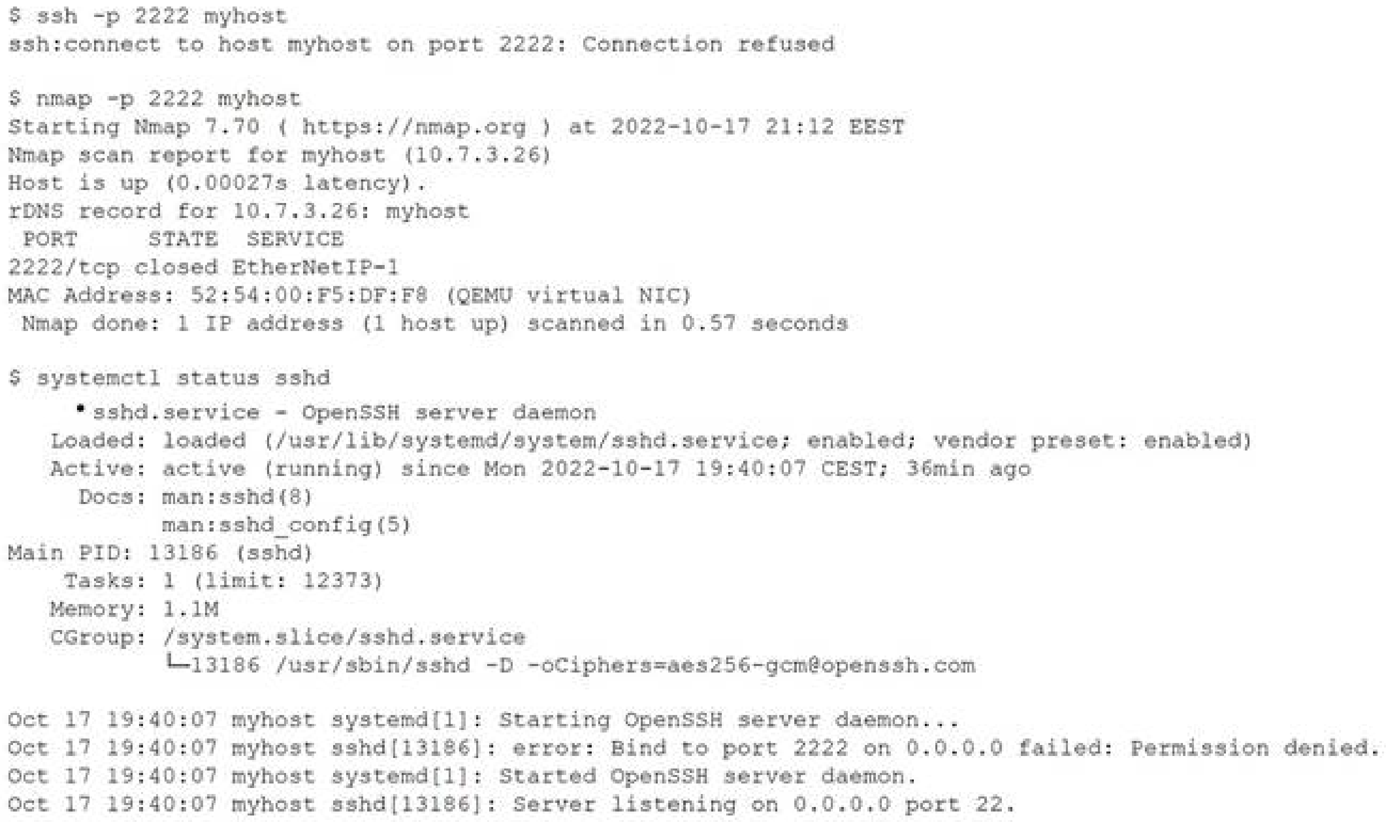

```
$ ssh -p 2222 myhost
ssh:connect to host myhost on port 2222: Connection refused

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

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

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

Which of the following commands will fix the issue?

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

Answer: A

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

NEW QUESTION 141
Which of the following enables administrators to configure and enforce MFA on a Linux system?

A. Kerberos
B. SELinux
C. PAM
D. PKI

Answer: C

Explanation:
The mechanism that enables administrators to configure and enforce MFA on a Linux system is PAM. PAM stands for Pluggable Authentication Modules, which 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 MFA, which stands for Multi-Factor Authentication, which is a security technique that requires the user to provide more than one piece of evidence to prove their identity. MFA can enhance the security of the system and prevent unauthorized access. PAM enables administrators to configure and enforce MFA on a Linux system. This is the correct answer to the question. The other options are incorrect because they either do not manage authentication and authorization on Linux systems (Kerberos or PKI) or do not support MFA (SELinux). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Implementing Basic Security, page 517.

NEW QUESTION 144
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 149
After installing some RPM packages, a systems administrator discovers the last package that was installed was not needed. Which of the following commands can be used to remove the package?

A. dnf remove packagename
B. apt-get remove packagename
C. rpm -i packagename
D. apt remove packagename

Answer: A

Explanation:
The command that can be used to remove an RPM package that was installed by mistake is dnf remove packagename. This command will use the DNF package manager to uninstall an RPM package and its dependencies from a Linux system that uses RPM-based distributions, such as Red Hat Enterprise Linux or CentOS. The DNF package manager handles dependency resolution and metadata searching for RPM packages.
The other options are not correct commands for removing an RPM package from a Linux system. The apt-get remove packagename and apt remove packagename commands are used to remove Debian packages from a Linux system that uses Debian-based distributions, such as Ubuntu or Debian. They are not compatible with RPM packages. The rpm -i packagename command is used to install an RPM package, not to remove it. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 9: Managing Software Packages; How to install/remove/query/update RPM packages in Linux (Cheat Sheet …

NEW QUESTION 154
The applications team is reporting issues when trying to access the web service hosted in a Linux system. The Linux systems administrator is reviewing the following outputs:
Output 1:
* httpd.service = The Apache HTTPD Server
Loaded: loaded (/usr/lib/systemd/system/httpd.service; disabled; vendor preset: disabled) Active: inactive (dead)
Docs: man:httpd(8) man:apachectl(8) Output 2:
16:51:16 up 28 min, 1 user, load average: 0.00, 0.00, 0.07
Which of the following statements best describe the root cause? (Select two).

A. The httpd service is currently started.
B. The httpd service is enabled to auto start at boot time, but it failed to start.
C. The httpd service was manually stopped.
D. The httpd service is not enabled to auto start at boot time.
E. The httpd service runs without problems.
F. The httpd service did not start during the last server reboot.

Answer: CD

Explanation:
The httpd.service is the Apache HTTPD Server, which is a web service that runs on Linux systems. The output 1 shows that the httpd.service is inactive (dead), which means that it is not running. The output 1 also shows that the httpd.service is disabled, which means that it is not enabled to auto start at boot time. Therefore, the statements C and D best describe the root cause of the issue. The statements A, B, E, and F are incorrect because they do not match the output 1. References: [How to Manage Systemd Services on a Linux System]

NEW QUESTION 158
A systems administrator has been tasked with disabling the nginx service from the environment to prevent it from being automatically and manually started. Which of the following commands will accomplish this task?

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

Answer: C

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

NEW QUESTION 159
A cloud engineer needs to remove all dangling images and delete all the images that do not have an associated container. Which of the following commands will help to accomplish this task?

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

Answer: A

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

NEW QUESTION 163
A cloud engineer wants to delete all unused networks that are not referenced by any container. Which of the following commands will achieve this goal?

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

Answer: C

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

NEW QUESTION 166
An administrator thinks that a package was installed using a snap. Which of the following commands can the administrator use to verify this information?

A. snap list
B. snap find
C. snap install
D. snap try

Answer: A

Explanation:
The snap list command is used to display the installed snaps on the system1. Snaps are self-contained software packages that can be installed and updated across different Linux distributions2. The snap list command shows the name, version, revision, developer and notes of each snap1.
The snap find command is used to search for snaps in the Snap Store, which is an online repository of snaps2. The snap install command is used to install snaps from the Snap Store or from a local file2. The snap try command is used to test a snap without installing it, by mounting a directory that contains the snap files2.
These commands are not useful for verifying if a package was installed using a snap.

NEW QUESTION 167
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 171
A Linux engineer finds multiple failed login entries in the security log file for application users. The Linux engineer performs a security audit and discovers a security issue. Given the following:

```
# grep -iE '*www*|db' /etc/passwd
www-data:x:502:502:www-data:/var/www:/bin/bash db:x: 505:505:db: /opt/db:/bin/bash
```

Which of the following commands would resolve the security issue?

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

Answer: D

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

NEW QUESTION 176
A systems engineer is adding a new 1GB XFS filesystem that should be temporarily mounted under /ops/app. Which of the following is the correct list of commands to achieve this goal?

A.

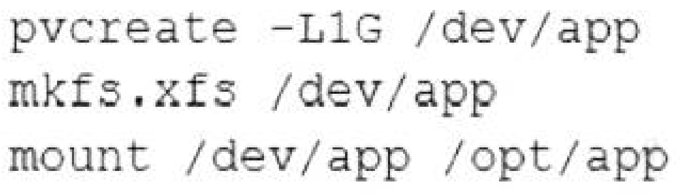

```
pvcreate -L1G /dev/app
mkfs.xfs /dev/app
mount /dev/app /opt/app
```

B.

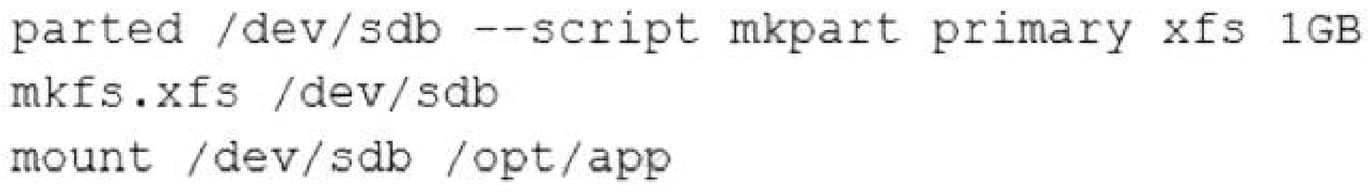

```
parted /dev/sdb --script mkpart primary xfs 1GB
mkfs.xfs /dev/sdb
mount /dev/sdb /opt/app
```

C.

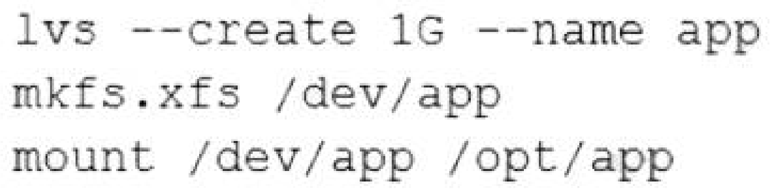

```
lvs --create 1G --name app
mkfs.xfs /dev/app
mount /dev/app /opt/app
```

D.

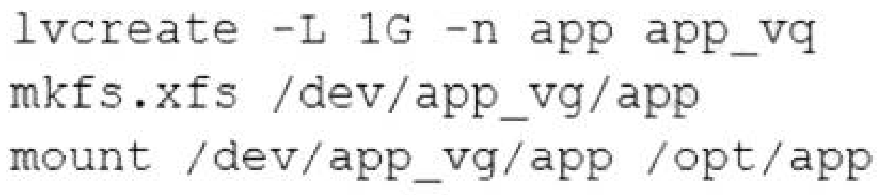

```
lvcreate -L 1G -n app app_vq
mkfs.xfs /dev/app_vg/app
mount /dev/app_vg/app /opt/app
```

Answer: D

Explanation:
The list of commands in option D is the correct way to achieve the goal. The commands are as follows:
? fallocate -l 1G /ops/app.img creates a 1GB file named app.img under the /ops directory.
? mkfs.xfs /ops/app.img formats the file as an XFS filesystem.
? mount -o loop /ops/app.img /ops/app mounts the file as a loop device under the /ops/app directory. The other options are incorrect because they either use the wrong commands (dd or truncate instead of fallocate), the wrong options (-t or - f instead of -o), or the wrong order of arguments (/ops/app.img /ops/app instead of /ops/app /ops/app.img). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 323-324.

NEW QUESTION 181
A systems administrator is notified that the mysqld process stopped unexpectedly. The systems administrator issues the following command: sudo grep –i -r ‘out of memory’ /var/log
The output of the command shows the following:
kernel: Out of memory: Kill process 9112 (mysqld) score 511 or sacrifice child.
Which of the following commands should the systems administrator execute NEXT to troubleshoot this issue? (Select two).

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

Answer: AF

Explanation:
The free -h command can be used to check the amount of free and used memory in the system in a human-readable format. This can help to troubleshoot the issue of mysqld being killed due to out of memory. The vmstat -a 1 4 command can be used to monitor the system’s virtual memory statistics, such as swap usage, paging activity, and memory faults, every one second for four times. This can help to identify any memory pressure or performance issues that may cause out of memory errors. The nc -v 127.0.0.1 3306 command would attempt to connect to the MySQL server on port 3306 and display any diagnostic messages, but this would not help to troubleshoot the memory issue. The renice -15 $(pidof mysql) command would change the priority of the mysql process to -15, but this

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

NEW QUESTION 186
A Linux engineer needs to create a custom script, cleanup.sh, to run at boot as part of the system services. Which of the following processes would accomplish this task?

A. Create a unit file in the /etc/default/ director
B. systemct1 enable cleanupsystemct1 is-enabled cleanup
C. Create a unit file in the /etc/ske1/ director
D. systemct1 enable cleanupsystemct1 is-enabled cleanup
E. Create a unit file in the /etc/systemd/system/ director
F. systemct1 enable cleanupsystemct1 is-enabled cleanup
G. Create a unit file in the /etc/sysctl.d/ director
H. systemct1 enable cleanupsystemct1 is-enabled cleanup

Answer: C

Explanation:
The process that will accomplish the task of creating a custom script to run at boot as part of the system services is:
? Create a unit file in the /etc/systemd/system/ directory. A unit file is a configuration
file that defines the properties and behavior of a systemd service. The systemd is a system and service manager that controls the startup and operation of Linux systems. The /etc/systemd/system/ directory is the location where the administrator can create and store custom unit files. The unit file should have a name that matches the name of the script, such as cleanup.service, and should contain the following sections and options:
? Run the command systemct1 enable cleanup. This command will enable the service and create the necessary symbolic links to start the service at boot.
? Run the command systemct1 is-enabled cleanup. This command will check the status of the service and confirm that it is enabled.
This process will create a custom script, cleanup.sh, to run at boot as part of the system services. This is the correct process to use to accomplish the task. The other options are incorrect because they either use the wrong directory for the unit file (/etc/default/, /etc/skel/, or /etc/sysctl.d/) or do not create a unit file at all. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing System Services, pages 457-459.

NEW QUESTION 188
A systems administrator created a new Docker image called test. After building the image, the administrator forgot to version the release. Which of the following will allow the administrator to assign the v1 version to the image?

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

Answer: C

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

NEW QUESTION 192
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:

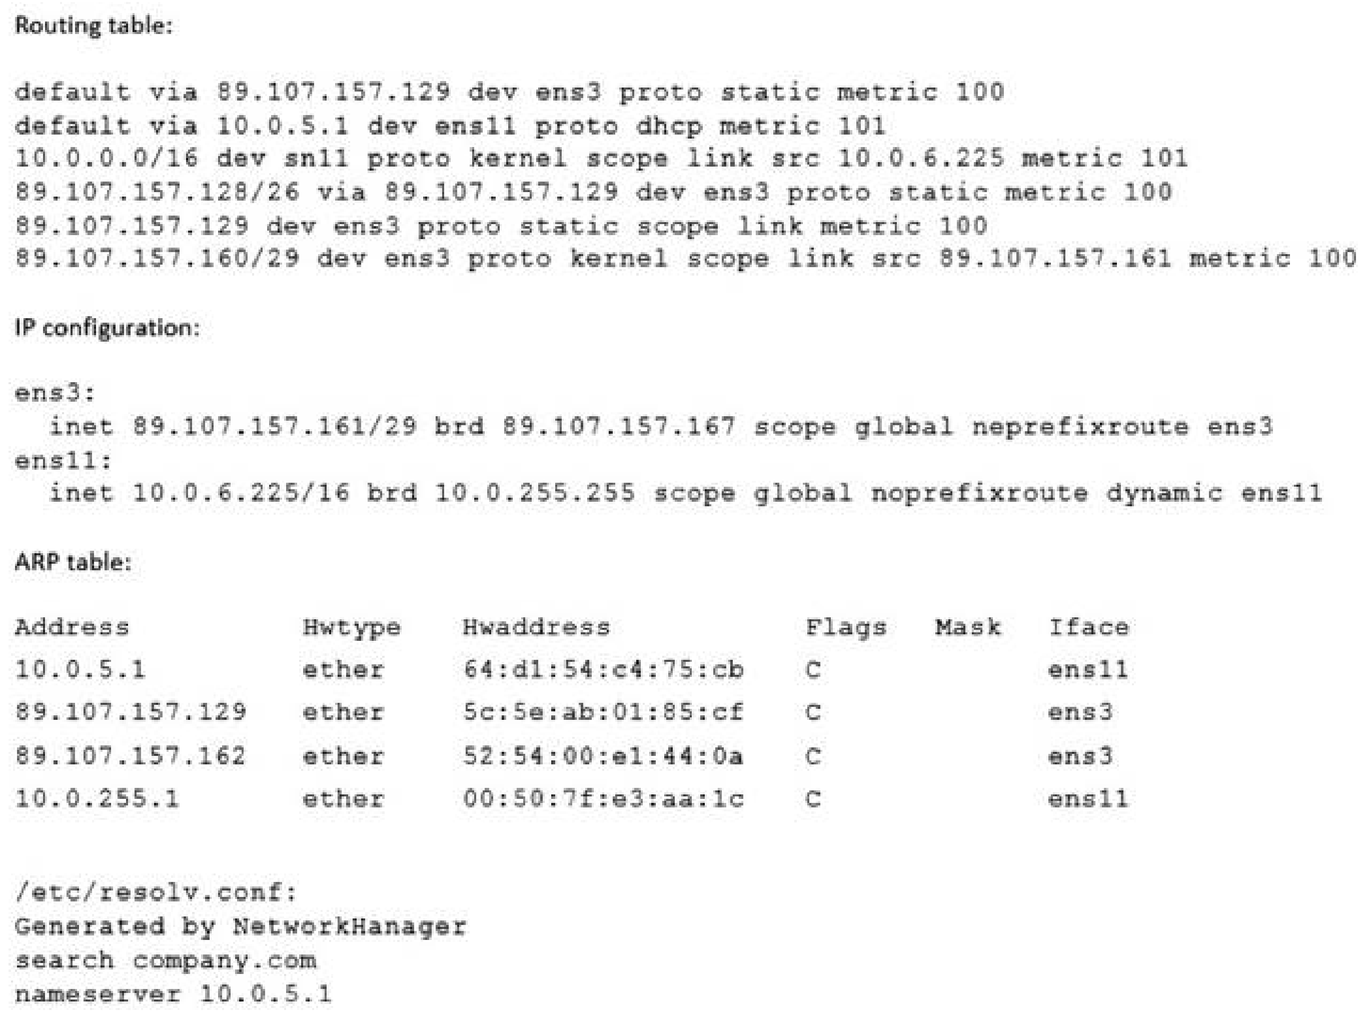

Routing table:

```
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 193
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 194
A Linux administrator created a new file system. Which of the following files must be updated to ensure the filesystem mounts at boot time?

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

Answer: C

Explanation:
The file that must be updated to ensure the filesystem mounts at boot time is /etc/fstab. This file contains information about the filesystems that are mounted

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

NEW QUESTION 196
A Linux administrator is reviewing changes to a configuration file that includes the following section:

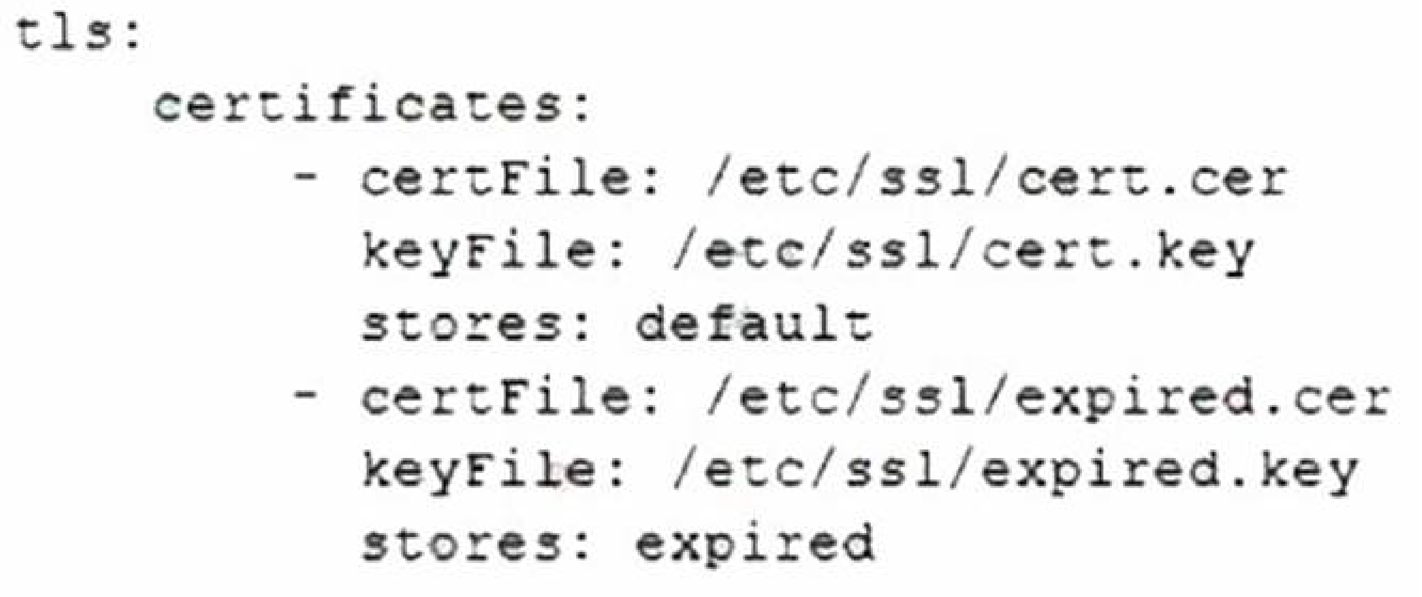

```
tls:
    certificates:
        - certFile: /etc/ssl/cert.cer
          keyFile: /etc/ssl/cert.key
          stores: default
        - certFile: /etc/ssl/expired.cer
          keyFile: /etc/ssl/expired.key
          stores: expired
```

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

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

Answer: C

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

NEW QUESTION 201
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 203
A Linux user reported the following error after trying to connect to the system remotely: ssh: connect to host 10.0.1.10 port 22: Resource temporarily unavailable
The Linux systems administrator executed the following commands in the Linux system while trying to diagnose this issue:

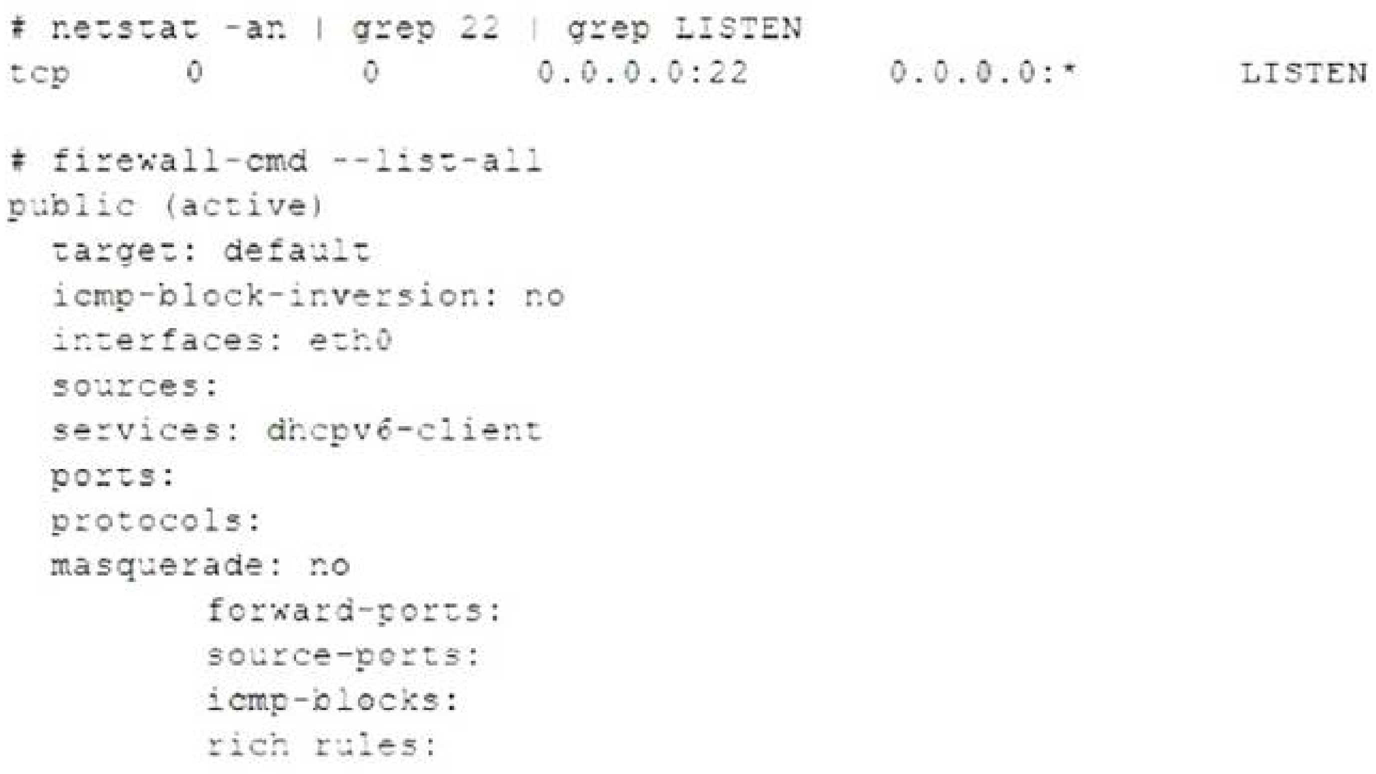

```
# netstat -an | grep 22 | grep LISTEN
tcp      0         0        0.0.0.0:22       0.0.0.0:*        LISTEN

# firewall-cmd --list-all
public (active)
  target: default
  icmp-block-inversion: no
  interfaces: eth0
  sources:
  services: dhcpv6-client
  ports:
  protocols:
  masquerade: no
         forward-ports:
         source-ports:
         icmp-blocks:
         rich rules:
```

Which of the following commands will resolve this issue?

A. firewall-cmd --zone=public --permanent --add-service=22
B. systemct1 enable firewalld; systemct1 restart firewalld
C. firewall-cmd --zone=public --permanent --add-service=ssh
D. firewall-cmd --zone=public --permanent --add-port=22/udp

Answer: C

Explanation:
The firewall-cmd --zone=public --permanent --add-service=ssh command will resolve the issue by allowing SSH connections on port 22 in the public zone of the firewalld service. This command will add the ssh service to the permanent configuration of the public zone, which means it will persist after a reboot or a reload of the firewalld service. The firewall-cmd --zone=public --permanent --add-service=22 command is invalid, as 22 is not a valid service name. The systemct1 enable firewalld; systemct1 restart firewalld command will enable and restart the firewalld service, but it will not change the firewall rules. The firewall-cmd --zone=public --permanent --add-port=22/udp command will allow UDP traffic on port 22 in the public zone, but SSH uses TCP, not UDP. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 543.

NEW QUESTION 207
A systems administrator is investigating a service that is not starting up. Given the following information:

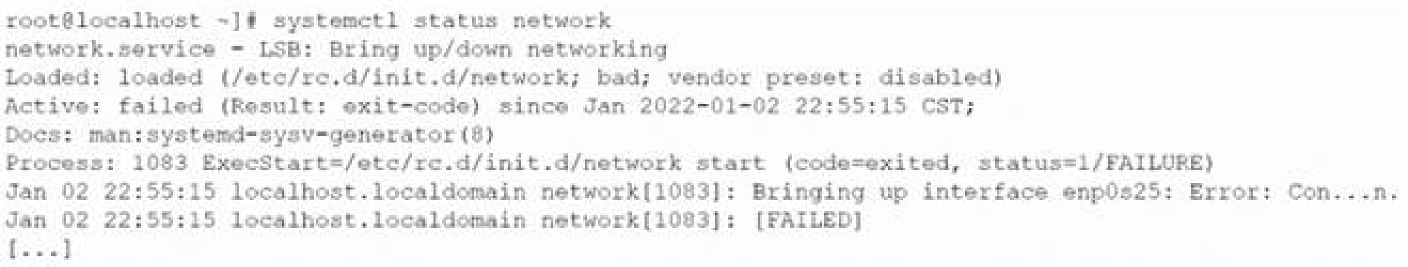

```
root@localhost ~]# systemctl status network
network.service - LSB: Bring up/down networking
Loaded: loaded (/etc/rc.d/init.d/network; bad; vendor preset: disabled)
Active: failed (Result: exit-code) since Jan 2022-01-02 22:55:15 CST;
Docs: man:systemd-sysv-generator(8)
Process: 1083 ExecStart=/etc/rc.d/init.d/network start (code=exited, status=1/FAILURE)
Jan 02 22:55:15 localhost.localdomain network[1083]: Bringing up interface enp0s25: Error: Con...n.
Jan 02 22:55:15 localhost.localdomain network[1083]: [FAILED]
[...]
```

Which of the following systemd commands should the administrator use in order to obtain more details about the failing service?

A. systemct1 analyze network
B. systemct1 info network
C. sysctl -a network
D. journalctl -xu network

Answer: D

Explanation:
The systemd is a system and service manager for Linux systems that provides a standard way to control and monitor system services. The systemd uses various commands and tools to manage and troubleshoot system services, such as systemct1, sysctl, and journalctl. The systemct1 command is used to start, stop, enable, disable, restart, reload, status, and list system services. The sysctl command is used to configure kernel parameters at runtime. The journalctl command is used to view and filter the logs of system services.
To investigate a service that is not starting up, the administrator can use the journalctl command with the -xu option. The -x option enables verbose output that includes explanatory text and priority information. The -u option filters the output by a specific unit name, such as network.service. Therefore, the command journalctl -xu network will show detailed logs of the network service, which can help identify the cause of the failure. The statement D is correct.
The statements A, B, and C are incorrect because they do not provide more details about the failing service. The systemct1 analyze network command does not exist.
The systemct1 info network command shows basic information about the network unit, such as description, load state, active state, sub state, and main PID. The sysctl -a network command shows all kernel parameters related to network settingsR. eferences: [How to Use Systemd to Manage System Services]

NEW QUESTION 212
......