Exam Questions XK0-005

CompTIA Linux+ Certification Exam

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

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

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

Answer: D

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

NEW QUESTION 2
A Linux user is trying to execute commands with sudo but is receiving the following error:
$ sudo visudo
>>> /etc/sudoers: syntax error near line 28 <<< sudo: parse error in /etc/sudoers near line 28 sudo: no valid sudoers sources found, quitting The following output is provided:
# grep root /etc/shadow root :* LOCK *: 14600 ::::::
Which of the following actions will resolve this issue?

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

Answer: B

NEW QUESTION 3
A user reported issues when trying to log in to a Linux server. The following outputs were received:
Given the outputs above. which of the following is the reason the user is una-ble to log in to the server?

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

Answer: D

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

NEW QUESTION 4
A Linux administrator wants to find out whether files from the wget package have been altered since they were installed. Which of the following commands will provide the correct information?

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

Answer: D

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

NEW QUESTION 5
In which of the following filesystems are system logs commonly stored?

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

Answer: A

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

NEW QUESTION 6
An administrator recently updated the BIND software package and would like to review the default configuration that shipped with this version. Which of the following files should the administrator review?

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

Answer: A

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

NEW QUESTION 7
A Linux administrator is alerted to a storage capacity issue on a server without a specific mount point or directory. Which of the following commands would be MOST helpful for troubleshooting? (Choose two.)

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

Answer: BD

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

NEW QUESTION 8
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 9
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 10
A Linux administrator is troubleshooting a systemd mount unit file that is not working correctly. The file contains:
[root@system] # cat mydocs.mount [Unit]
Description=Mount point for My Documents drive [Mount]
What=/dev/drv/disk/by-uuid/94afc9b2-ac34-ccff-88ae-297ab3c7ff34 Where=/home/user1/My Documents
Options=defaults Type=xfs
[Install]
WantedBy=multi-user.target
The administrator verifies the drive UUID correct, and user1 confirms the drive should be
mounted as My Documents in the home directory. Which of the following can the administrator
do to fix the issues with mounting the drive? (Select two).

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

Answer: AE

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

NEW QUESTION 10
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 15
During a security scan, the password of an SSH key file appeared to be too weak and was cracked. Which of the following commands would allow a user to choose a stronger password and set it on the existing SSH key file?

A. passwd
B. ssh
C. ssh-keygen

D. pwgen

Answer: C

Explanation:
The command that would allow a user to choose a stronger password and set it on the existing SSH key file is ssh-keygen -p -f <keyfile>. This command uses the ssh-keygen tool, which is used to generate, manage, and convert authentication keys for SSH. The -p option stands for passphrase, and it allows the user to change or remove the passphrase of an existing private key file. The -f option specifies the filename of the key file. The command will prompt the user for the old passphrase, and then for the new passphrase twice.
The other options are not correct commands for changing the password of an SSH key file. The passwd command is used to change the password of a user account on a Linux system, not an SSH key file. The ssh command is used to log in to a remote system using SSH, not to change the password of an SSH key file. The pwgen command is used to generate random passwords, not to change the password of an SSH key file.
References: ssh-keygen(1) - Linux manual page; How To: Change Passphrase for SSH Private Key - Unix Tutorial

NEW QUESTION 16
A Linux administrator is troubleshooting a memory-related issue. Based on the output of the commands:

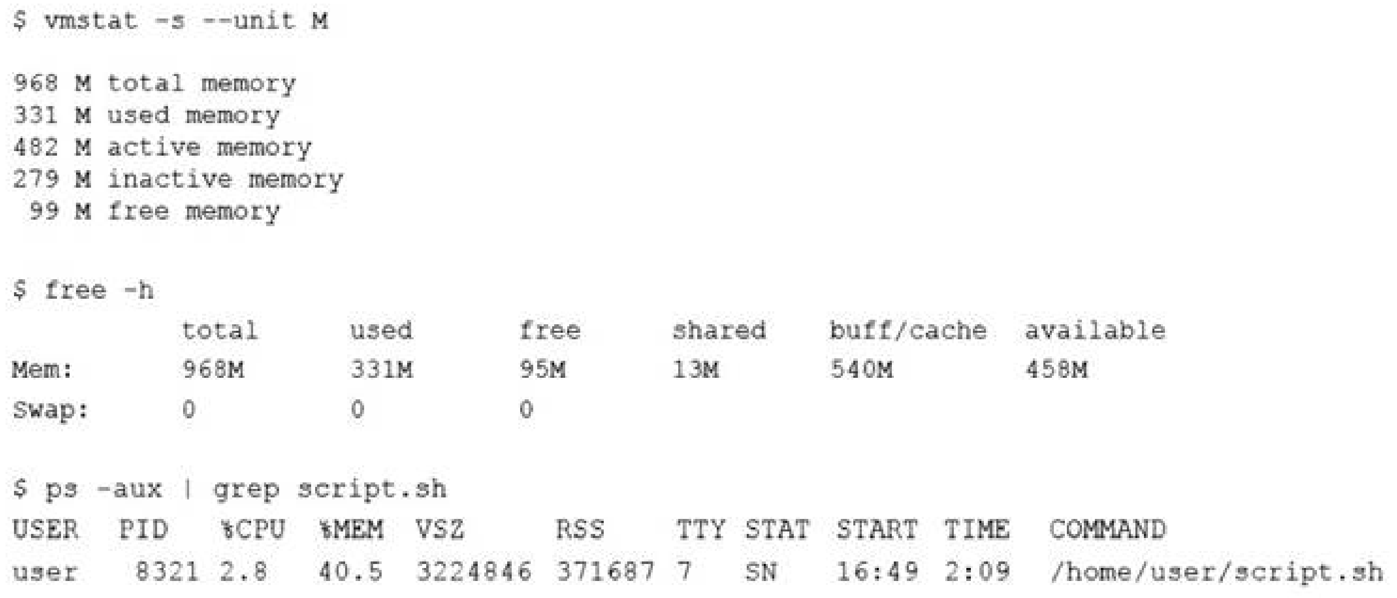

```
$ vmstat -s --unit M

968 M total memory
331 M used memory
482 M active memory
279 M inactive memory
 99 M free memory


$ free -h
          total      used       free      shared     buff/cache   available
Mem:      968M       331M       95M       13M        540M         458M
Swap:     0          0          0


$ ps -aux | grep script.sh
USER   PID    %CPU  %MEM  VSZ      RSS     TTY  STAT  START  TIME   COMMAND
user    8321 2.8    40.5  3224846  371687  7    SN    16:49  2:09   /home/user/script.sh
```

Which of the following commands would address the issue?

A. top -p 8321
B. kill -9 8321
C. renice -10 8321
D. free 8321

Answer: B

Explanation:
The command that would address the memory-related issue is kill -9 8321. This command will send a SIGKILL signal to the process with the PID 8321, which is the mysqld process that is using 99.7% of the available memory according to the top output. The SIGKILL signal will terminate the process immediately and free up the memory it was using. However, this command should be used with caution as it may cause data loss or corruption if the process was performing some critical operations.
The other options are not correct commands for addressing the memory-related issue. The top -p 8321 command will only display information about the process with the PID 8321, but will not kill it or reduce its memory usage. The renice -10 8321 command will change the priority (niceness) of the process with the PID 8321 to -10, which means it will have a higher scheduling priority, but this will not affect its memory consumption. The free 8321 command is invalid because free does not take a PID as an argument; free only displays information about the total, used, and free memory in the system. References: How to troubleshoot Linux server memory issues; kill(1) - Linux manual page

NEW QUESTION 20
Some servers in an organization have been compromised. Users are unable to access to the organization’s web page and other services. While reviewing the system log, a systems administrator notices messages from the kernel regarding firewall rules:

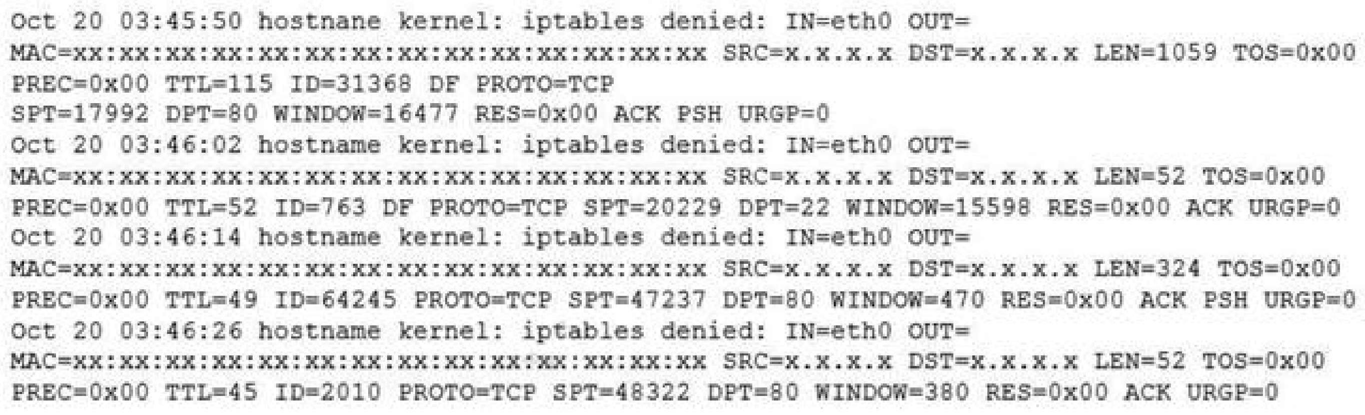

```
Oct 20 03:45:50 hostnane kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx SRC=x.x.x.x DST=x.x.x.x LEN=1059 TOS=0x00
PREC=0x00 TTL=115 ID=31368 DF PROTO=TCP
SPT=17992 DPT=80 WINDOW=16477 RES=0x00 ACK PSH URGP=0
Oct 20 03:46:02 hostname kernel: iptables denied: IN=eth0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

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

A.

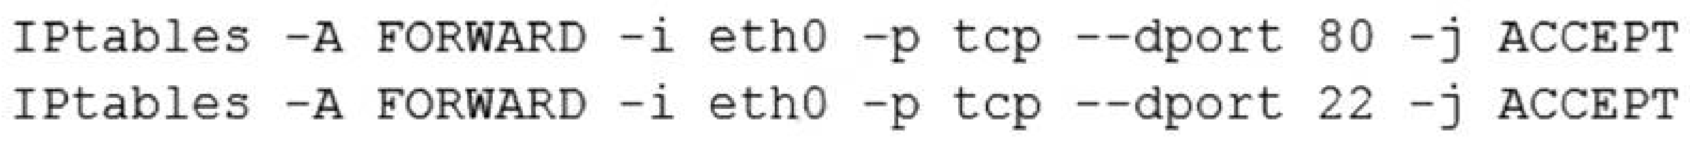

```
IPtables -A FORWARD -i eth0 -p tcp --dport 80 -j ACCEPT
IPtables -A FORWARD -i eth0 -p tcp --dport 22 -j ACCEPT
```

B.

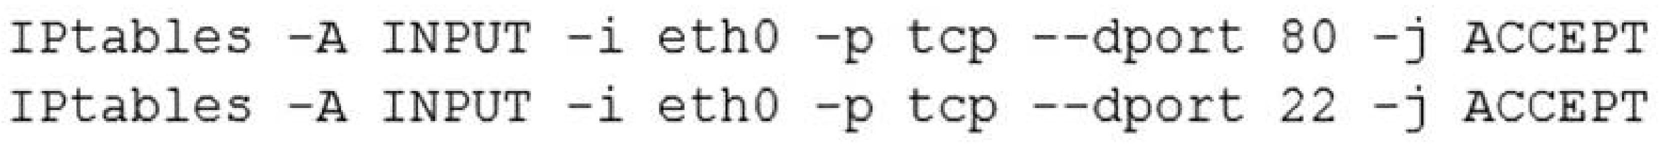

```
IPtables -A INPUT -i eth0 -p tcp --dport 80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --dport 22 -j ACCEPT
```

C.

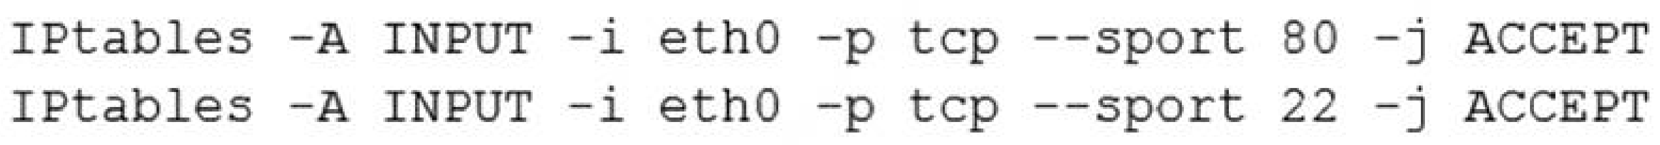

```
IPtables -A INPUT -i eth0 -p tcp --sport 80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --sport 22 -j ACCEPT
```

D.

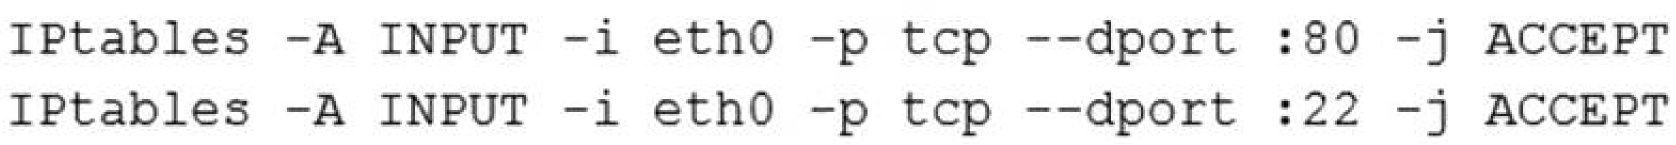

```
IPtables -A INPUT -i eth0 -p tcp --dport :80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --dport :22 -j ACCEPT
```

Answer: A

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

NEW QUESTION 23
User1 is a member of the accounting group. Members of this group need to be able to execute but not make changes to a script maintained by User2. The script should not be accessible to other users or groups. Which of the following will give proper access to the script?

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

Answer: A

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

NEW QUESTION 25
Rugged appliances are small appliances with ruggedized hardware and like Quantum Spark appliance they use which operating system?

A. Centos Linux
B. Gaia embedded
C. Gaia
D. Red Hat Enterprise Linux version 5

Answer: B

Explanation:
Rugged appliances are small appliances with ruggedized hardware that use Gaia embedded as their operating system. Gaia embedded is a version of Gaia that is optimized for embedded devices such as Rugged appliances and Quantum Spark appliances. Gaia embedded supports features such as VPN, firewall, identity awareness, application control, URL filtering, and anti-bot. Gaia embedded does not use Centos Linux, Gaia, or Red Hat Enterprise Linux version 5 as their operating system. References: Check Point Rugged Appliance Datasheet, page 1.

NEW QUESTION 30
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 35
A user generated a pair of private-public keys on a workstation. Which of the following commands will allow the user to upload the public key to a remote server and enable passwordless login?

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

Answer: D

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

NEW QUESTION 38
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 42
A Linux administrator wants to prevent the httpd web service from being started both manually and automatically on a server. Which of the following should the administrator use to accomplish this task?

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

Answer: A

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

NEW QUESTION 45
Users are reporting that writes on a system configured with SSD drives have been taking longer than expected, but reads do not seem to be affected. A Linux systems administrator is investigating this issue and working on a solution. Which of the following should the administrator do to help solve the issue?

A. Run the corresponding command to trim the SSD drives.
B. Use fsck on the filesystem hosted on the SSD drives.

C. Migrate to high-density SSD drives for increased performance.
D. Reduce the amount of files on the SSD drives.

Answer: A

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

NEW QUESTION 47
A user is unable to remotely log on to a server using the server name server1 and port 22.
The Linux engineer troubleshoots the issue and gathers the following information: Which of the following is most likely causing the issue?

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

Answer: B

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

NEW QUESTION 48
A Linux administrator is installing a web server and needs to check whether web traffic has already been allowed through the firewall. Which of the following commands should the administrator use to accomplish this task?

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

Answer: C

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

NEW QUESTION 51
A systems administrator wants to be sure the sudo rules just added to /etc/sudoers are valid. Which of the following commands can be used for this task?

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

Answer: A

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

NEW QUESTION 54
A Linux systems administrator is troubleshooting an I/O latency on a single CPU server. The administrator runs a top command and receives the following output:
%Cpu(s): 0.2 us, 33.1 sy, 0.0 ni, 0.0 id, 52.4 wa, 0.0 hi, 0.2 si, 0.0 st
Which of the following is correct based on the output received from the exe-cuted command?

A. The server's CPU is taking too long to process users' requests.
B. The server's CPU shows a high idle-time value.

C. The server's CPU is spending too much time waiting for data inputs.
D. The server's CPU value for the time spent on system processes is low.

Answer: C

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

NEW QUESTION 59
A systems administrator is gathering information about a file type and the contents of a file. Which of the following commands should the administrator use to accomplish this task?

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

Answer: A

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

NEW QUESTION 62
An engineer needs to insert a character at the end of the current line in the vi text editor. Which of the following will allow the engineer to complete this task?

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

Answer: D

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

NEW QUESTION 63
A DevOps engineer wants to allow the same Kubernetes container configurations to be deployed in development, testing, and production environments. A key requirement is that the containers should be configured so that developers do not have to statically configure custom, environment-specific locations. Which of the following should the engineer use to meet this requirement?

A. Custom scheduler
B. Node affinity
C. Overlay network
D. Ambassador container

Answer: D

Explanation:
To allow the same Kubernetes container configurations to be deployed in different environments without statically configuring custom locations, the engineer can use an ambassador container (D). An ambassador container is a proxy container that handles communication between containers and external services. It can dynamically configure locations based on environment variables or other methods. The other options are not related to this requirement. References:
? [CompTIA Linux+ Study Guide], Chapter 11: Working with Containers, Section: Using Ambassador Containers
? [How to Use Ambassador Containers]

NEW QUESTION 65
A developer has been unable to remove a particular data folder that a team no longer uses. The developer escalated the issue to the systems administrator. The following output was received:

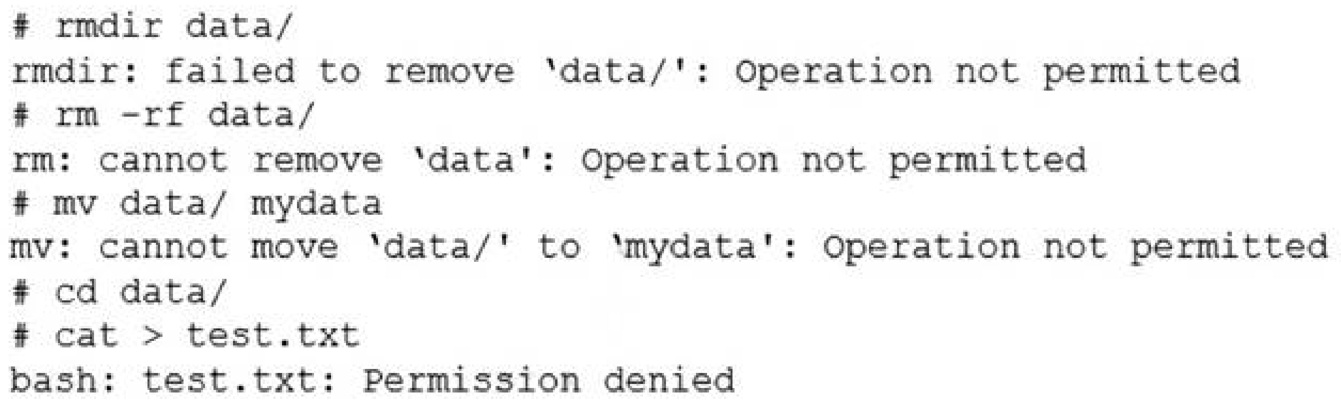

```
# rmdir data/
rmdir: failed to remove 'data/': Operation not permitted
# rm -rf data/
rm: cannot remove 'data': Operation not permitted
# mv data/ mydata
mv: cannot move 'data/' to 'mydata': Operation not permitted
# cd data/
# cat > test.txt
bash: test.txt: Permission denied
```

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

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

Answer: C

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

NEW QUESTION 66
A systems administrator wants to delete app . conf from a Git repository. Which of the following commands will delete the file?

A. git tag ap
B. conf
C. git commit app . conf
D. git checkout app . conf
E. git rm ap
F. conf

Answer: D

Explanation:
To delete a file from a Git repository, the administrator can use the command git rm app.conf (D). This will remove the file “app.conf” from the working directory and stage it for deletion from the repository. The administrator can then commit the change with git commit -m "Delete app.conf" to finalize the deletion. The other commands will not delete the file, but either tag, commit, or checkout the file. References:
? [CompTIA Linux+ Study Guide], Chapter 10: Working with Git, Section: Deleting Files with Git
? [How to Delete Files from Git]

NEW QUESTION 70
A Linux systems administrator receives reports from various users that an application hosted on a server has stopped responding at similar times for several days in a row. The administrator logs in to the system and obtains the following output:
Output 1:

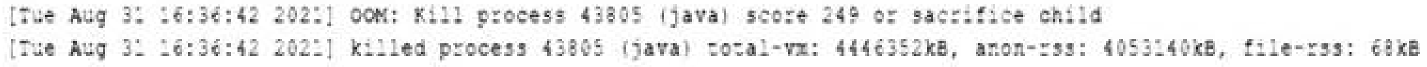

```
[Tue Aug 31 16:36:42 2021] OOM: Kill process 43805 (java) score 249 or sacrifice child
[Tue Aug 31 16:36:42 2021] killed process 43805 (java) total-vm: 4446352kB, anon-rss: 4053140kB, file-rss: 68kB
```

Output 2:

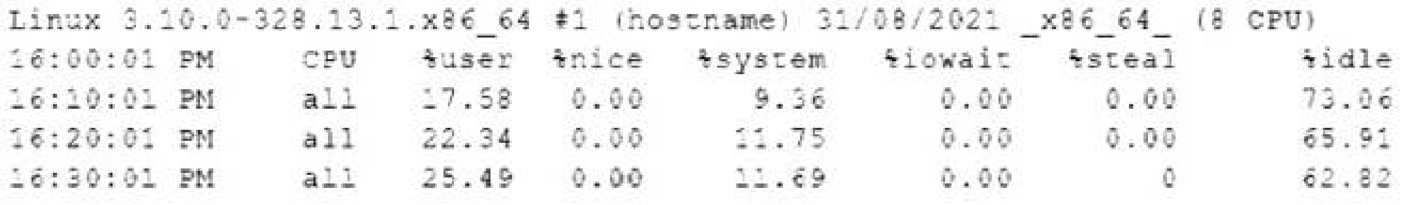

```
Linux 3.10.0-328.13.1.x86_64 #1 (hostname) 31/08/2021 _x86_64_ (8 CPU)
16:00:01 PM     CPU    %user   %nice   %system   %iowait   %steal      %idle
16:10:01 PM     all    17.58    0.00      9.36      0.00     0.00      73.06
16:20:01 PM     all    22.34    0.00     11.75      0.00     0.00      65.91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

Output 3:

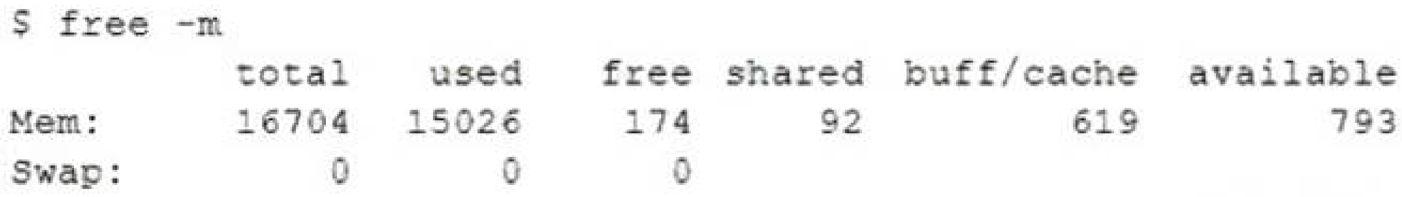

```
$ free -m
          total    used    free  shared  buff/cache   available
Mem:      16704   15026     174      92         619         793
Swap:         0       0       0
```

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

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

D. Configure the swap space to allow for spikes in usage during peak hours and prevent the Java process from stopping due to a lack of memory.

Answer: D

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

NEW QUESTION 73
A Linux system fails to start and delivers the following error message:

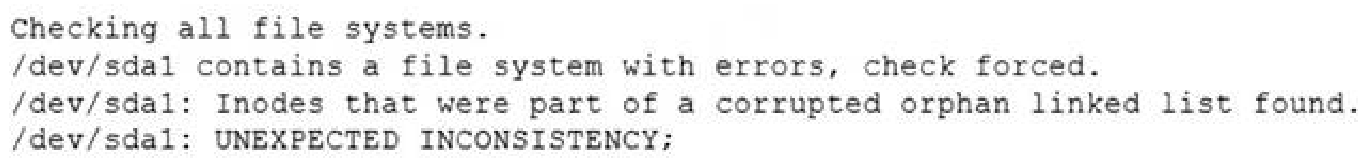

```
Checking all file systems.
/dev/sda1 contains a file system with errors, check forced.
/dev/sda1: Inodes that were part of a corrupted orphan linked list found.
/dev/sda1: UNEXPECTED INCONSISTENCY;
```

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 75
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 77
Users in the human resources department are trying to access files in a newly created directory. Which of the following commands will allow the users access to the files?

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

Answer: B

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

? The web search result 2 explains how to use the chgrp command with examples.
? The web search result 3 compares the chmod and chgrp commands and their effects on file permissions.

NEW QUESTION 80
A systems administrator is enabling LUKS on a USB storage device with an ext4 filesystem format. The administrator runs dmesg and notices the following output:

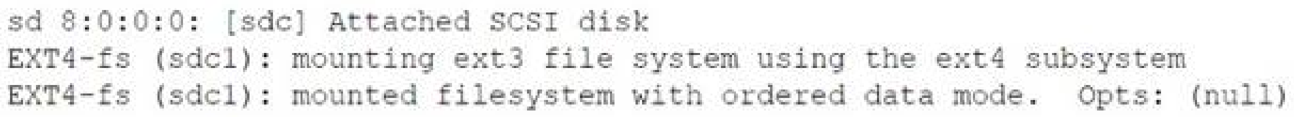

```
sd 8:0:0:0: [sdc] Attached SCSI disk
EXT4-fs (sdc1): mounting ext3 file system using the ext4 subsystem
EXT4-fs (sdc1): mounted filesystem with ordered data mode.  Opts: (null)
```

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

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

Answer: CDH

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

NEW QUESTION 84
A Linux administrator copied a Git repository locally, created a feature branch, and committed some changes to the feature branch. Which of the following Git actions should the Linux administrator use to publish the changes to the main branch of the remote repository?

A. rebase
B. tag
C. commit
D. push

Answer: D

Explanation:
The push action is used to publish the changes made in a local branch to a remote branch of a Git repository. This action will update the remote branch with the commits made in the local branch and synchronize the two branches. The rebase action is used to reapply commits from one branch onto another branch, creating a linear history of commits. This action does not publish any changes to a remote repository. The tag action is used to create an annotated reference to a specific commit in a Git repository. This action does not publish any changes to a remote repository. The commit action is used to record changes made in the local repository and create a new snapshot of the project state. This action does not publish any changes to a remote repository. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 20: Writing and Executing Bash Shell Scripts, page 579.

NEW QUESTION 89
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 90
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 95
A newly created container has been unable to start properly, and a Linux administrator is analyzing the cause of the failure. Which of the following will allow the administrator to determine the FIRST command that is executed inside the container right after it starts?

A. docker export <container_id>
B. docker info <container_id>
C. docker start <container_id>
D. docker inspect <container_id>

Answer: D

Explanation:
The command that will allow the administrator to determine the first command that is executed inside the container right after it starts is docker inspect <container_id>. This command will display detailed information about the container, including its configuration, state, network settings, mounts, and logs. One of the configuration fields is “Entrypoint”, which shows the command that is executed when the container is run. The entrypoint can be specified in the Dockerfile or overridden at runtime using the --entrypoint option.
The other options are not correct commands for determining the first command that is executed inside the container. The docker export <container_id> command will export the contents of the container’s filesystem as a tar archive to STDOUT. This will not show the entrypoint of the container, but only its files. The docker info <container_id> command is invalid because docker info does not take any arguments. It shows system-wide information about Docker, such as the number of containers, images, volumes, networks, and storage drivers. The docker start <container_id> command will start a stopped container and attach its STDOUT and STDERR to the terminal. This will not show the entrypoint of the container, but only its output. References: docker inspect | Docker Docs; docker export | Docker Docs; docker info | Docker Docs; docker start | Docker Docs

NEW QUESTION 98
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.

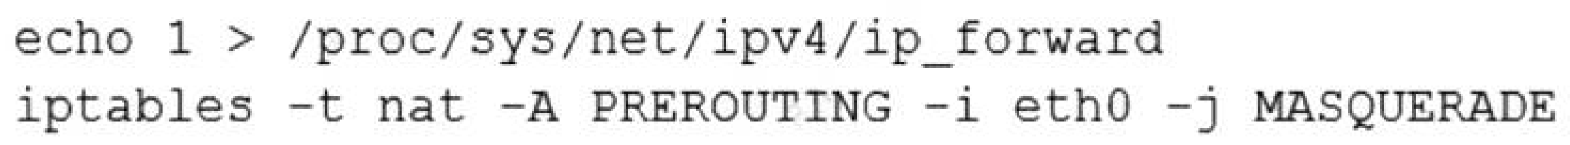

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A PREROUTING -i eth0 -j MASQUERADE
```

A.

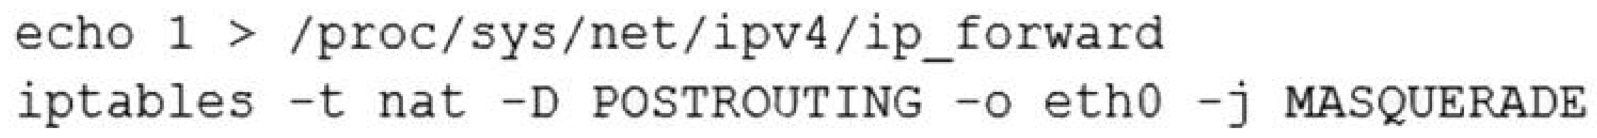

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -D POSTROUTING -o eth0 -j MASQUERADE
```

B.

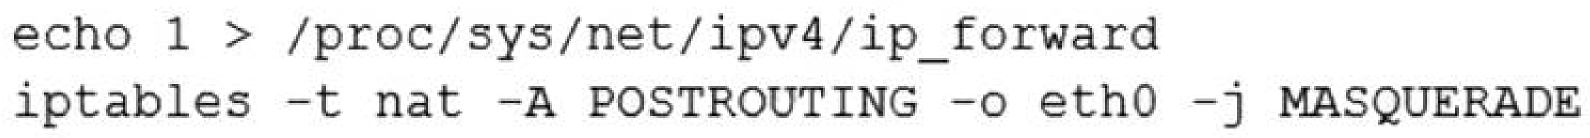

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
```

C.

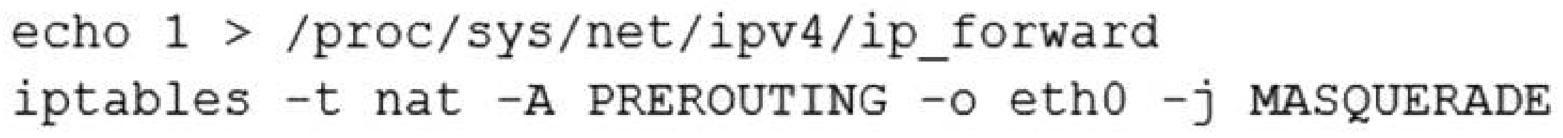

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A PREROUTING -o eth0 -j MASQUERADE
```

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 101
A systems administrator is tasked with preventing logins from accounts other than root, while the file /etc/nologin exists. Which of the following PAM modules will accomplish this task?

A. pam_login.so

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

Answer: D

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

NEW QUESTION 104
Following the migration from a disaster recovery site, a systems administrator wants a server to require a user to change credentials at initial login. Which of the following commands should
be used to ensure the aging attribute?

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

Answer: B

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

NEW QUESTION 107
A systems administrator checked out the code from the repository, created a new branch, made changes to the code, and then updated the main branch. The systems administrator wants to ensure that the Terraform state files do not appear in the main branch. Which of following should the administrator use to meet this requirement?

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

Answer: B

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

NEW QUESTION 108
An administrator is trying to diagnose a performance issue and is reviewing the following output:

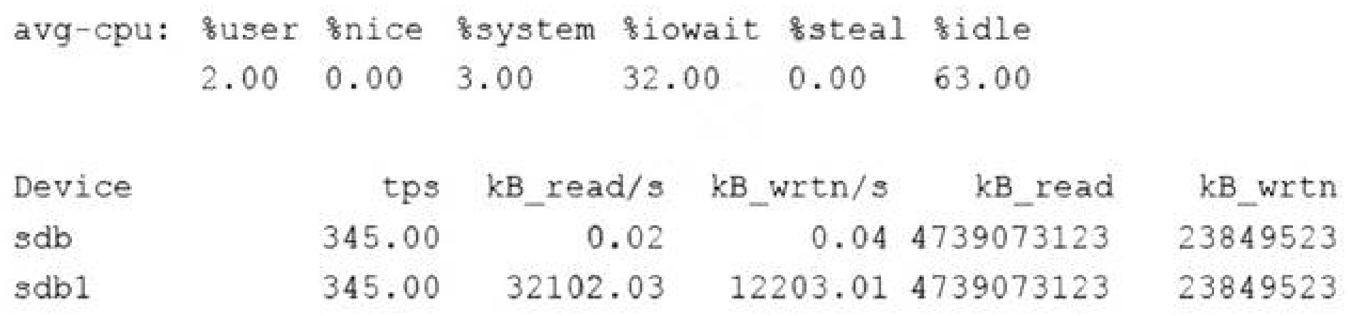

```
avg-cpu: %user  %nice  %system  %iowait  %steal  %idle
          2.00   0.00    3.00    32.00     0.00   63.00

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

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

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

Answer: B

Explanation:
The system has reached its maximum permitted throughput, therefore iowait
is increasing. The output of iostat -x shows that the device sda has an average throughput of 44.01 MB/s, which is equal to the disk maximum throughput of 44 Mbps. The output also shows that the device sda has an average iowait of 99.99%, which means that the CPU is waiting for the disk to complete the I/O requests. This indicates that the disk is the bottleneck and the system is slow due to the high iowait. The other options are incorrect because they are not supported by the outputs. The system has not reached its maximum IOPS, as the device sda has an average IOPS of 563.50, which is lower than the disk maximum IOPS of 690.

The system is not mostly idle, as the output of top shows that the CPU is 100% busy. The system does not have a partitioned disk, as the output of lsblk shows that the device sda has only one partition sda1. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Optimizing Linux Systems, pages 513-514.

NEW QUESTION 110

A systems administrator needs to check if the service systemd-resolved.service is running without any errors. Which of the following commands will show this information?

A. systemct1 status systemd-resolved.service
B. systemct1 enable systemd-resolved.service
C. systemct1 mask systemd-resolved.service
D. systemct1 show systemd-resolved.service

Answer: A

Explanation:

The command systemct1 status systemd-resolved.service will show the information about the service systemd-resolved.service. The systemct1 command is a tool for managing system services and units. The status option displays the current status of a unit, such as active, inactive, or failed. The output also shows the unit description, loaded configuration, process ID, memory usage, and recent log messages. This command will show if the service systemd-resolved.service is running without any errors. This is the
correct command to use to accomplish the task. The other options are incorrect because they either perform different actions (enable, mask, or show) or do not show the status of the service (systemct1 show systemd-resolved.service only shows the properties of the service, not the status). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, page 427.

NEW QUESTION 114

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

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

Answer: A

Explanation:

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

NEW QUESTION 119

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

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

Answer: B

Explanation:

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

NEW QUESTION 123

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

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

Answer: D

Explanation:

Terraform is a tool for building, changing, and managing infrastructure as code in a cloud- based environment. Terraform uses configuration files to describe the

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

NEW QUESTION 126
A Linux systems administrator is setting up a new web server and getting 404 - NOT FOUND errors while trying to access the web server pages from the browser. While working on the diagnosis of this issue, the Linux systems administrator executes the following commands:

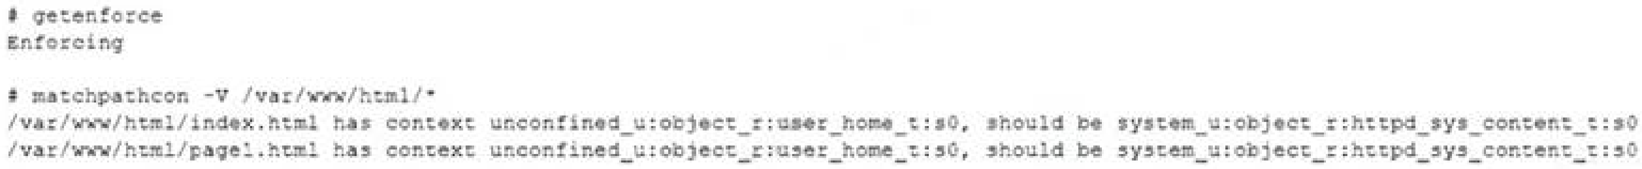

```
# getenforce
Enforcing

# matchpathcon -V /var/www/html/*
/var/www/html/index.html has context unconfined_u:object_r:user_home_t:s0, should be system_u:object_r:httpd_sys_content_t:s0
/var/www/html/page1.html has context unconfined_u:object_r:user_home_t:s0, should be system_u:object_r:httpd_sys_content_t:s0
```

Which of the following commands will BEST resolve this issue?

A. sed -i 's/SELINUX=enforcing/SELINUX=disabled/' /etc/selinux/config
B. restorecon -R -v /var/www/html
C. setenforce 0
D. setsebool -P httpd_can_network_connect_db on

Answer: B

Explanation:
The command restorecon -R -v /var/www/html will best resolve the issue. The issue is caused by the incorrect SELinux context of the web server files under th /var/www/html directory. The output of ls -Z /var/www/html shows that the files have the type user_home_t, which is not allowed for web content. The command restorecon restores the default SELinux context of files based on the policy rules. The options -R and -v are used to apply the command recursively and verbosely. This command will change the type
of the files to httpd_sys_content_t, which is the correct type for web content. This will allow the web server to access the files and serve the pages to the browser. The other options are incorrect because they either disable SELinux entirely (sed -i 's/SELINUX=enforcing/SELINUX=disabled/' /etc/selinux/config or setenforce 0), which is not a good security practice, or enable an unnecessary boolean (setsebool -P httpd_can_network_connect_db on), which is not related to the issue.
References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 535.

NEW QUESTION 127
DRAG DROP
A new drive was recently added to a Linux system. Using the environment and tokens provided, complete the following tasks:
• Create an appropriate device label.
• Format and create an ext4 file system on the new partition. The current working directory is /.

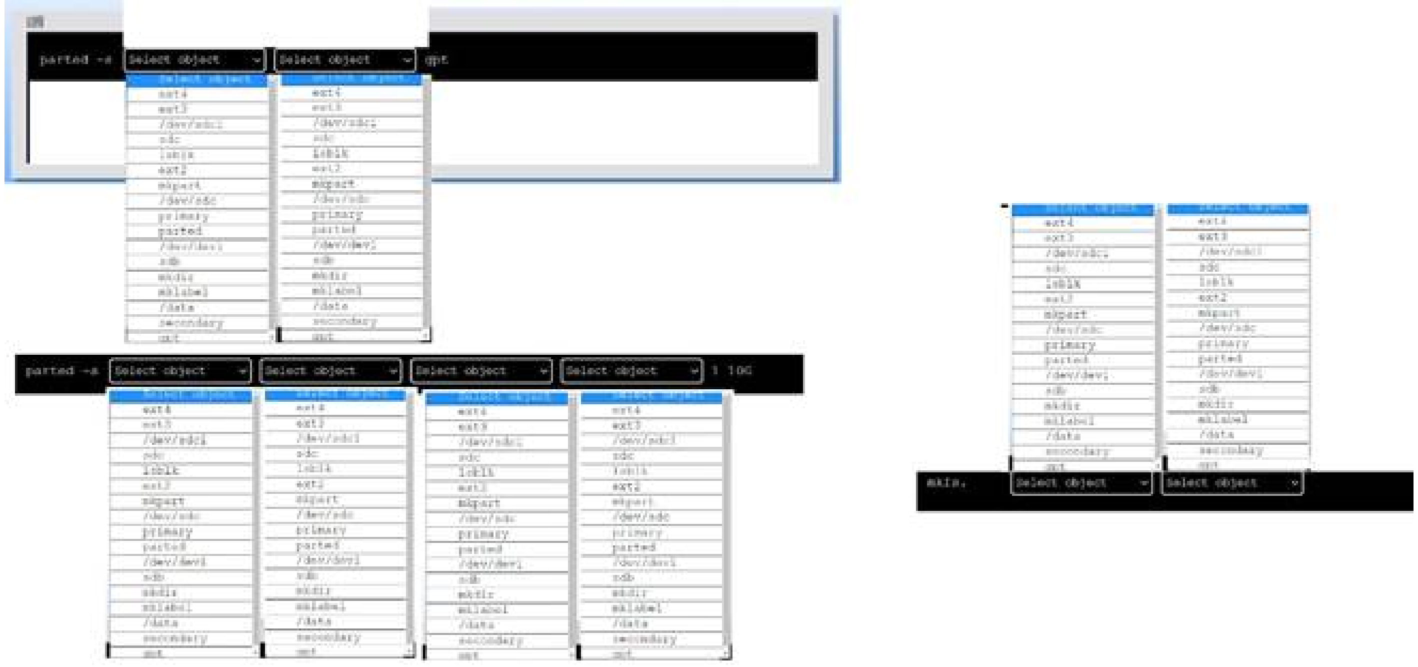

```
parted -s [Select object] [Select object] gpt
parted -s [Select object] [Select object] [Select object] [Select object] 1 10G
mkfs. [Select object] [Select object]
```

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 131
An administrator attempts to connect to a remote server by running the following command:
$ nmap 192.168.10.36
Starting Nmap 7.60 (https://nmap.org) at 2022-03-29 20:20 UTC Nmap scan report for www1 (192.168.10.36)
Host is up (0.000091s latency). Not shown: 979 closed ports PORT STATE SERVICE 21/tcp open ftp 22/tcp filtered ssh 631/tcp open ipp
Nmap done: 1 IP address (1 host up) scanned in 0.06 seconds
Which of the following can be said about the remote server?

A. A firewall is blocking access to the SSH server.
B. The SSH server is not running on the remote server.
C. The remote SSH server is using SSH protocol version 1.
D. The SSH host key on the remote server has expired.

Answer: A

Explanation:
This is because the port 22/tcp is shown as filtered by nmap, which means that nmap cannot determine whether the port is open or closed because a firewall or other device is blocking its probes. If the SSH server was not running on the remote server, the port would be shown as closed, which means that nmap received a TCP RST packet in response to its probe. If the remote SSH server was using SSH protocol version 1, the port would be shown as open, which means that nmap received a TCP SYN/ACK packet in response to its probe. If the SSH host key on the remote server had expired, the port would also be
shown as open, but the SSH client would display a warning message about the host key verification failure. Therefore, the best explanation for the filtered state of the port 22/tcp is that a firewall is preventing nmap from reaching the SSH server.
You can find more information about nmap port states and how to interpret them in the following web search results:
? Nmap scan what does STATE=filtered mean?
? How to find ports marked as filtered by nmap
? Technical Tip: NMAP scan shows ports as filtered

NEW QUESTION 133
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:

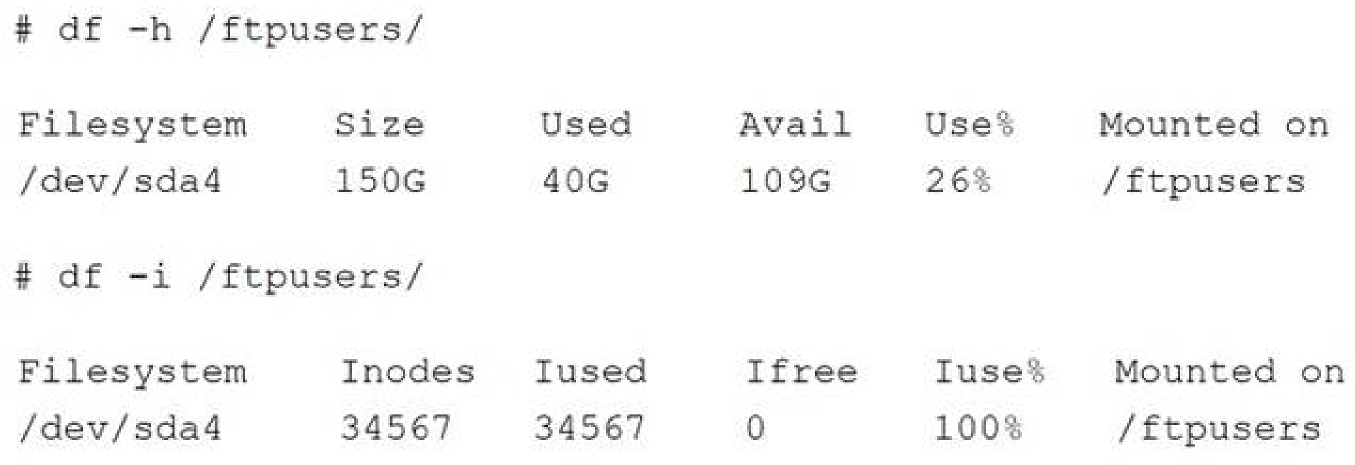

```
# df -h /ftpusers/

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 136
An administrator deployed a Linux server that is running a web application on port 6379/tcp.
SELinux is in enforcing mode based on organization policies. The port is open on the firewall.
Users who are trying to connect to a local instance of the web application receive Error 13, Permission denied.
The administrator ran some commands that resulted in the following output:

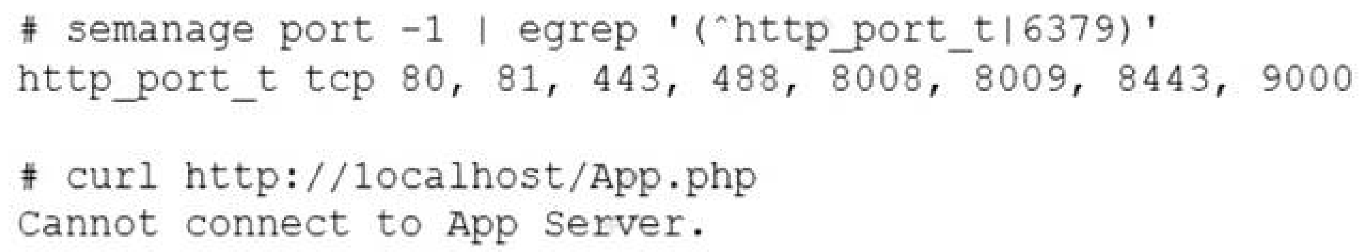

```
# semanage port -l | egrep '(^http_port_t|6379)'
http_port_t tcp 80, 81, 443, 488, 8008, 8009, 8443, 9000

# curl http://localhost/App.php
Cannot connect to App Server.
```

Which of the following commands should be used to resolve the issue?

A. semanage port -d -t http_port_t -p tcp 6379
B. semanage port -a -t http_port_t -p tcp 6379
C. semanage port -a http_port_t -p top 6379
D. semanage port -l -t http_port_tcp 6379

Answer: B

Explanation:
The command semanage port -a -t http_port_t -p tcp 6379 adds a new port definition to the SELinux policy and assigns the type http_port_t to the port 6379/tcp. This allows the web application to run on this port and accept connections from users. This is the correct way to resolve the issue. The other options are incorrect because they either delete a port definition (-d), use the wrong protocol (top instead of tcp), or list the existing port definitions (-l). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 535.

NEW QUESTION 137
A systems administrator is encountering performance issues. The administrator runs 3 commands with the following output

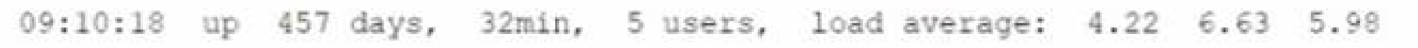

```
09:10:18  up  457 days,  32min,  5 users,  load average:  4.22  6.63  5.98
```

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

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

Answer: A

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

NEW QUESTION 138
A new Linux systems administrator just generated a pair of SSH keys that should allow connection to the servers. Which of the following commands can be used to copy a key file to remote servers? (Choose two.)

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

Answer: DF

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

NEW QUESTION 139
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 142
A Linux administrator is troubleshooting an issue in which an application service failed to start on a Linux server. The administrator runs a few commands and gets the following outputs:

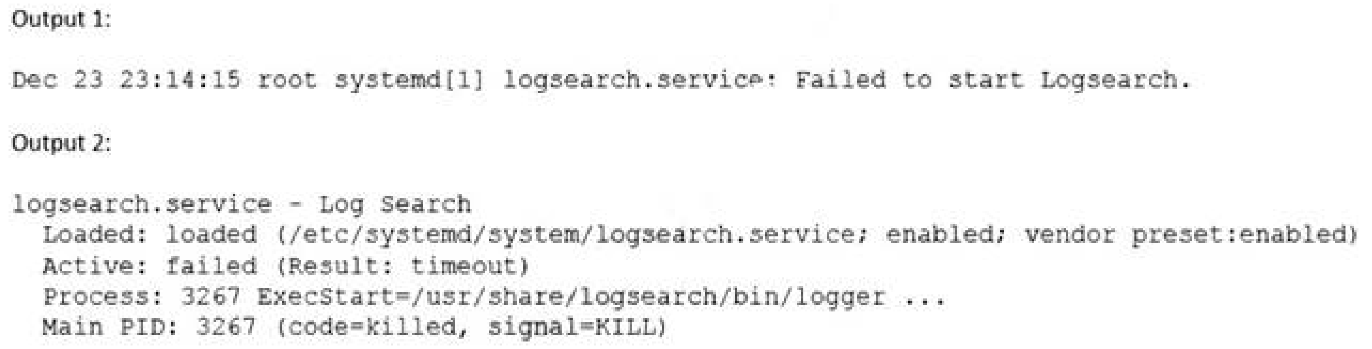

Output 1:

```
Dec 23 23:14:15 root systemd[1] logsearch.service: Failed to start Logsearch.
```

Output 2:

```
logsearch.service - Log Search
  Loaded: loaded (/etc/systemd/system/logsearch.service; enabled; vendor preset:enabled)
  Active: failed (Result: timeout)
  Process: 3267 ExecStart=/usr/share/logsearch/bin/logger ...
  Main PID: 3267 (code=killed, signal=KILL)
```

Based on the above outputs, which of the following is the MOST likely action the administrator should take to resolve this issue?

A. Enable the logsearch.service and restart the service.
B. Increase the TimeoutStartUSec configuration for the logsearch.sevice.
C. Update the OnCalendar configuration to schedule the start of the logsearch.service.
D. Update the KillSignal configuration for the logsearch.service to use TERM.

Answer: B

Explanation:
The administrator should increase the TimeoutStartUSec configuration for the logsearch.service to resolve the issue. The output of systemct1 status logsearch.service shows that the service failed to start due to a timeout. The output of cat /etc/systemd/system/logsearch.service shows that the service has a TimeoutStartUSec configuration of 10 seconds, which might be too short for the service to start. The administrator should increase this value to a higher number, such as 30 seconds or 1 minute, and then restart the service. The other options are incorrect because they are not related to the issue. The service is already enabled, as shown by the output of systemct1 is-enabled logsearch.service. The service does not use an OnCalendar configuration, as it is not a timer unit. The service does not use a KillSignal configuration, as it is not being killed by a signal. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, pages 434-435.

NEW QUESTION 147
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 150
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 152
A Linux administrator reviews a set of log output files and needs to identify files that contain any occurrence of the word denied. All log files containing entries in uppercase or lowercase letters should be included in the list. Which of the following commands should the administrator use to accomplish this task?

A. find . -type f -print | xrags grep -ln denied
B. find . -type f -print | xrags grep -nv denied
C. find . -type f -print | xrags grep -wL denied
D. find . -type f -print | xrags grep -li denied

Answer: D

Explanation:
The command find . -type f -print | xargs grep -li denied will accomplish the task of identifying files that contain any occurrence of the word denied. The find command is a tool for searching for files and directories on Linux systems. The . is the starting point of the search, which means the current directory. The -type f option specifies the type of the file, which means regular file. The -print option prints the full file name on the standard output. The | is a pipe symbol that redirects the output of one command to the input of another command. The xargs command is a tool for building and executing commands from standard input. The grep command is a tool for searching for patterns in files or input.
The -li option specifies the flags that the grep command should apply. The -l flag shows only the file names that match the pattern, instead of the matching lines. The -i flag ignores the case of the pattern, which means it matches both uppercase and lowercase letters.
The denied is the pattern that the grep command should search for. The command find . - type f -print | xargs grep -li denied will find all the regular files in the current directory and its subdirectories, and then search for any occurrence of the word denied in those files, ignoring the case, and print only the file names that match the pattern. This will allow the administrator to identify files that contain any occurrence of the word denied. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not ignore the case of the pattern (find . -type f -print | xargs grep -ln denied or find . -type f -print | xargs grep -wL denied) or do not show the file names that match the pattern (find . -type f -print | xargs grep -nv denied). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Managing Logging and Monitoring, page 489.

NEW QUESTION 156
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 159
An administrator accidentally deleted the /boot/vmlinuz file and must resolve the issue before the server is rebooted. Which of the following commands should the administrator use to identify the correct
version of this file?

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

Answer: A

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

NEW QUESTION 163
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 166
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:

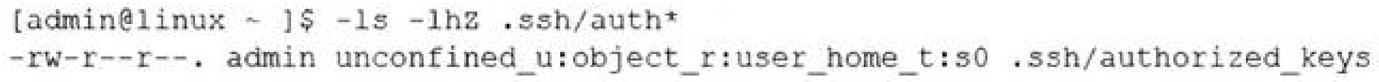

```
[admin@linux ~ ]$ -ls -lhZ .ssh/auth*
-rw-r--r--. admin unconfined_u:object_r:user_home_t:s0 .ssh/authorized_keys
```

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 169
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 171
Developers have requested implementation of a persistent, static route on the application server. Packets sent over the interface eth0 to 10.0.213.5/32 should be routed via 10.0.5.1. Which of the following commands should the administrator run to achieve this goal?

A. route -i etho -p add 10.0.213.5 10.0.5.1
B. route modify eth0 +ipv4.routes "10.0.213.5/32 10.0.5.1"
C. echo "10.0.213.5 10.0.5.1 eth0" > /proc/net/route
D. ip route add 10.0.213.5/32 via 10.0.5.1 dev eth0

Answer: D

Explanation:
The command ip route add 10.0.213.5/32 via 10.0.5.1 dev eth0 adds a static route to the routing table that sends packets destined for 10.0.213.5/32 (a single host) through the gateway 10.0.5.1 on the interface eth0. This is the correct way to achieve the goal. The other options are incorrect because they either use the

wrong syntax (route -i etho -p add), the wrong command (route modify), or the wrong file (/proc/net/route). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 379.

NEW QUESTION 173
A systems administrator is tasked with creating a cloud-based server with a public IP address.

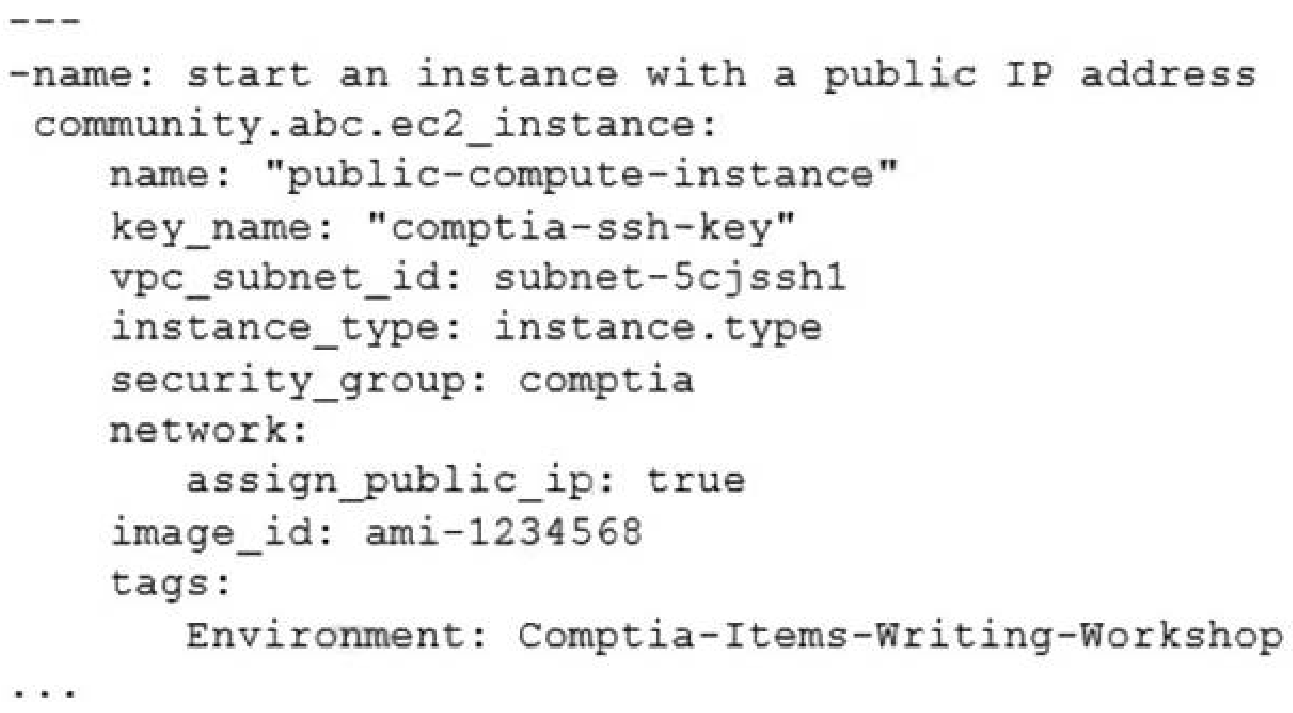

```
---
-name: start an instance with a public IP address
 community.abc.ec2_instance:
    name: "public-compute-instance"
    key_name: "comptia-ssh-key"
    vpc_subnet_id: subnet-5cjssh1
    instance_type: instance.type
    security_group: comptia
    network:
       assign_public_ip: true
    image_id: 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 177
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 180
A systems administrator is checking the system logs. The administrator wants to look at the last 20 lines of a log. Which of the following will execute the command?

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

Answer: B

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

NEW QUESTION 182
A Linux administrator needs to ensure that Java 7 and Java 8 are both locally available for developers to use when deploying containers. Currently only Java 8 is available. Which of the following commands should the administrator run to ensure both versions are available?

A. docker image load java:7
B. docker image pull java:7
C. docker image import java:7
D. docker image build java:7

Answer: B

Explanation:
The command that the administrator should run to ensure that both Java 7 and Java 8 are locally available for developers to use when deploying containers is docker image pull java:7. This command will use the docker image pull subcommand to download the java:7 image from Docker Hub, which is the default registry for Docker images. The java:7 image contains Java 7 installed on a Debian-based Linux system. The administrator can also specify a different registry by using the syntax registry/repository:tag.
The other options are not correct commands for ensuring that both Java 7 and Java 8 are locally available for developers to use when deploying containers. The docker image load java:7 command will load an image from a tar archive or STDIN, not from a registry. The docker image import java:7 command will create a new filesystem image from the contents of a tarball, not from a registry. The docker image build java:7 command will build an image from a Dockerfile, not from a registry. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks; docker image pull | Docker Docs

NEW QUESTION 185
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 188
A systems administrator configured firewall rules using firewalld. However, after the system is rebooted, the firewall rules are not present:

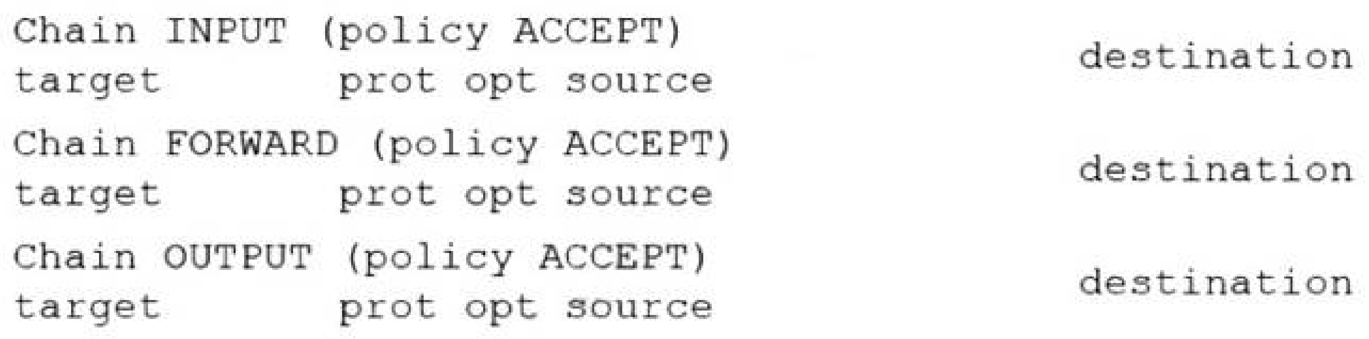

```
Chain INPUT (policy ACCEPT)
target        prot opt source                  destination
Chain FORWARD (policy ACCEPT)
target        prot opt source                  destination
Chain OUTPUT (policy ACCEPT)
target        prot opt source                  destination
```

The systems administrator makes additional checks:

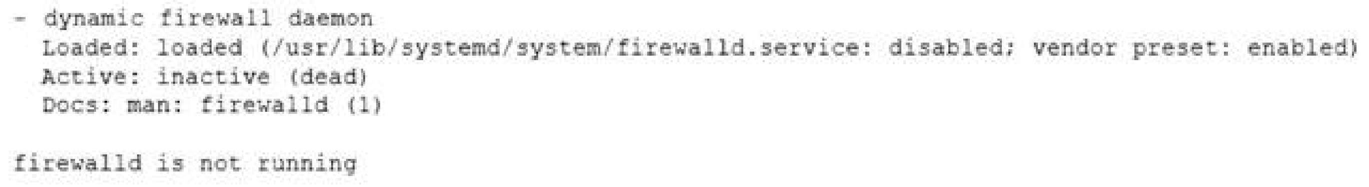

```
- dynamic firewall daemon
  Loaded: loaded (/usr/lib/systemd/system/firewalld.service: disabled; vendor preset: enabled)
  Active: inactive (dead)
  Docs: man: firewalld (1)

firewalld is not running
```

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

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

Answer: D

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

NEW QUESTION 192
A Linux engineer receives reports that files created within a certain group are being modified by users who are not group members. The engineer wants to reconfigure the server so that only file owners and group members can modify new files by default. Which of the following commands would accomplish this task?

A. chmod 775
B. umas
C. 002
D. chactr -Rv
E. chown -cf

Answer: B

Explanation:
The command umask 002 will accomplish the task of reconfiguring the server so that only file owners and group members can modify new files by default.
The umask command is a tool for setting the default permissions for new files and directories on Linux systems. The umask value is a four-digit octal number that represents the permissions that are subtracted from the default permissions. The default permissions for files are 666, which means read and write for owner, group, and others. The default permissions for directories are 777, which means read, write, and execute for owner, group, and others. The umask value consists of four digits: the first digit is for special permissions, such as setuid, setgid, and sticky bit; the second digit is for the owner permissions; the third digit is for the group permissions; and the fourth digit is for the others permissions. The umask value can be calculated by subtracting the desired permissions from the default permissions. For example, if the desired permissions for files are 664, which means read and write for owner and group, and read for others, then the umask value is 002, which is 666 - 664. The command umask 002 will set the umask value to 002, which will ensure that only file owners and group members can modify new files by default. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not set the default permissions for new files (chmod 775 or chown - cf) or do not exist (chattr -Rv). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing File Permissions and Ownership, page 349.

NEW QUESTION 193
A Linux administrator needs to resolve a service that has failed to start. The administrator runs the following command:

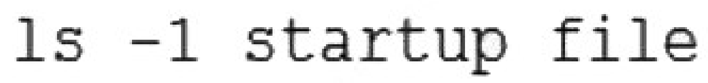

```
ls -1 startup file
```

The following output is returned

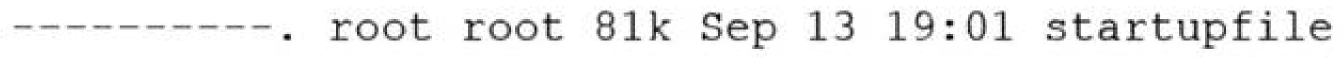

```
----------. root root 81k Sep 13 19:01 startupfile
```

Which of the following is MOST likely the issue?

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

Answer: A

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

NEW QUESTION 195
A systems administrator intends to use a UI-JID to mount a new partition per-manently on a Linux system. Which of the following commands can the adminis-trator run to obtain information about the UUlDs of all disks attached to a Linux system?

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

Answer: B

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

NEW QUESTION 197
A Linux administrator modified the SSH configuration file. Which of the following commands should be used to apply the configuration changes?

A. systemct1 stop sshd
B. systemct1 mask sshd
C. systemct1 reload sshd
D. systemct1 start sshd

Answer: C

Explanation:
The systemct1 reload sshd command can be used to apply the configuration changes of the SSH server daemon without restarting it. This is useful to avoid interrupting existing connections. The systemct1 stop sshd command would stop the SSH server daemon, not apply the changes. The systemct1 mask sshd command would prevent the SSH server daemon from being started, not apply the changes. The systemct1 start sshd command would start the SSH server daemon if it is not running, but it would not apply the changes if it is already running. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Secure Shell (SSH), page 415.

NEW QUESTION 198
A Linux administrator needs to correct the permissions of a log file on the server. Which of the following commands should be used to set filename.log permissions to -rwxr—r--. ?

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

Answer: A

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

NEW QUESTION 201
A systems administrator wants to test the route between IP address 10.0.2.15 and IP address 192.168.1.40. Which of the following commands will accomplish this task?

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

Answer: B

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

NEW QUESTION 203
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 206
Users are experiencing high latency when accessing a web application served by a Linux machine. A systems administrator checks the network interface counters and sees the following:

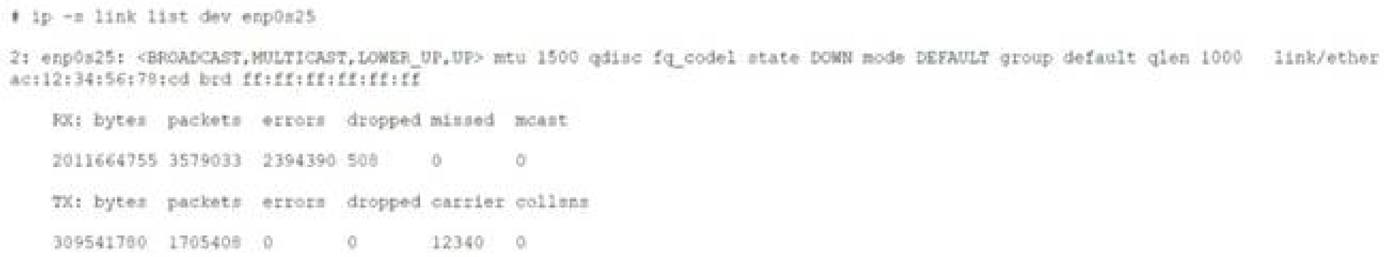

```
# ip -s link list dev enp0s25

2: enp0s25: <BROADCAST,MULTICAST,LOWER_UP,UP> mtu 1500 qdisc fq_codel state DOWN mode DEFAULT group default qlen 1000   link/ether
ac:12:34:56:78:cd brd ff:ff:ff:ff:ff:ff

    RX: bytes  packets  errors  dropped missed  mcast

    2011664755 3579033  2394390 508     0       0

    TX: bytes  packets  errors  dropped carrier collsns

    309541780  1705408  0       0       12340   0
```

Which of the following is the most probable cause of the observed latency?

A. The network interface is disconnected.
B. A connection problem exists on the network interface.
C. No IP address is assigned to the interface.
D. The gateway is unreachable.

Answer: B

Explanation:
The high number of errors and dropped packets in the output of the network interface counters indicate a connection problem on the network interface.
References:
? CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Networking, Section: Troubleshooting Network Issues, Page 359.
? Linux+ (Plus) Certification, Exam Objectives: 4.3 Given a scenario, troubleshoot and resolve basic network configuration and connectivity issues.

NEW QUESTION 208
A new file was added to a main Git repository. An administrator wants to synchronize a local copy with the contents of the main repository. Which of the following commands should the administrator use for this task?

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

Answer: B

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

NEW QUESTION 213
A Linux administrator needs to analyze a failing application that is running inside a container. Which of the following commands allows the Linux administrator to enter the running container and analyze the logs that are stored inside?

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

Answer: B

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

NEW QUESTION 214
Ann, a security administrator, is performing home directory audits on a Linux server. Ann issues the su Joe command and then issues the Is command. The output displays files that reside in Ann's home directory instead of Joe's. Which of the following represents the command Ann should have issued in order to list Joe's files?

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

Answer: A

Explanation:
The su command is used to switch to another user account on Linux systems. The - option makes the shell a login shell, which means that it will read the profile and environment variables of the target user. Without this option, the shell will retain the environment variables of the original user. This can cause confusion when

issuing commands that depend on these variables, such as ls, which uses the $HOME variable to determine the home directory. Therefore, Ann should have issued su - Joe to list Joe’s files instead of her own. References: [How to Use su Command in Linux with Examples]

NEW QUESTION 219
A user is attempting to log in to a Linux server that has Kerberos SSO ena-bled. Which of the following commands should the user run to authenticate and then show the ticket grants? (Select TWO).

A. kinit
B. klist
C. kexec
D. kioad
E. pkexec
F. realm

Answer: AB

Explanation:
The following commands can help the user to authenticate and show the ticket grants using Kerberos SSO on a Linux server:
? kinit: This command obtains and caches an initial ticket-granting ticket (TGT) for
the user from the Kerberos key distribution center (KDC). The user needs to enter their password or use a keytab file to authenticate1.
? klist: This command lists the cached tickets, including the TGT and any service tickets, for the user. It also shows the expiration time and flags for each ticket2.
For example, the user can run the following commands to log in and view their tickets:
$ kinit username@REALM Password for username@REALM:
$ klist
Ticket cache: FILE:/tmp/krb5cc_1000 Default principal: username@REALM
Valid starting Expires Service principal
04/06/2023 16:06:59 04/07/2023 02:06:59 krbtgt/REALM@REALM
renew until 04/13/2023 16:06:59 References:
? kinit(1) - Linux man page, section “Description”.
? klist(1) - Linux man page, section “Description”.

NEW QUESTION 221
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.
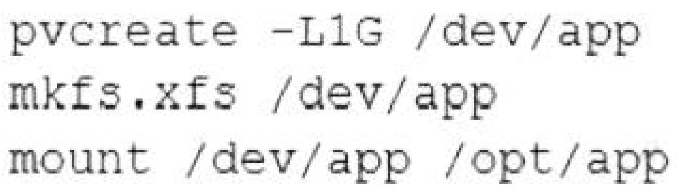

```
pvcreate -L1G /dev/app
mkfs.xfs /dev/app
mount /dev/app /opt/app
```

B.
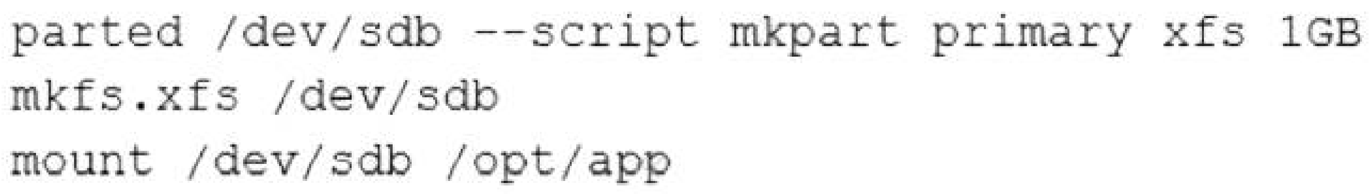

```
parted /dev/sdb --script mkpart primary xfs 1GB
mkfs.xfs /dev/sdb
mount /dev/sdb /opt/app
```

C.
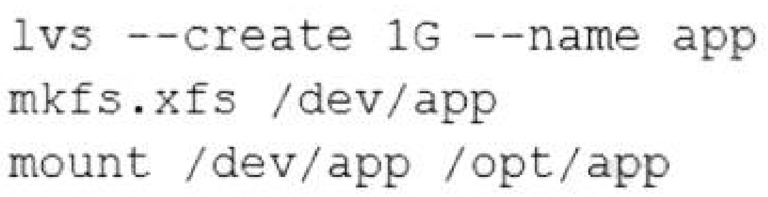

```
lvs --create 1G --name app
mkfs.xfs /dev/app
mount /dev/app /opt/app
```

D.
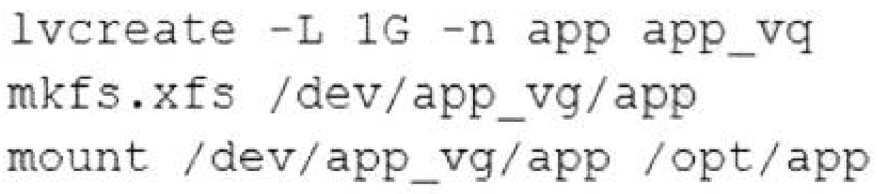

```
lvcreate -L 1G -n app app_vq
mkfs.xfs /dev/app_vg/app
mount /dev/app_vg/app /opt/app
```

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 223
A Linux administrator was tasked with deleting all files and directories with names that are contained in the sobelete.txt file. Which of the following commands will accomplish this task?

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

Answer: D

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

NEW QUESTION 226
A systems administrator detected corruption in the /data filesystem. Given the following output:

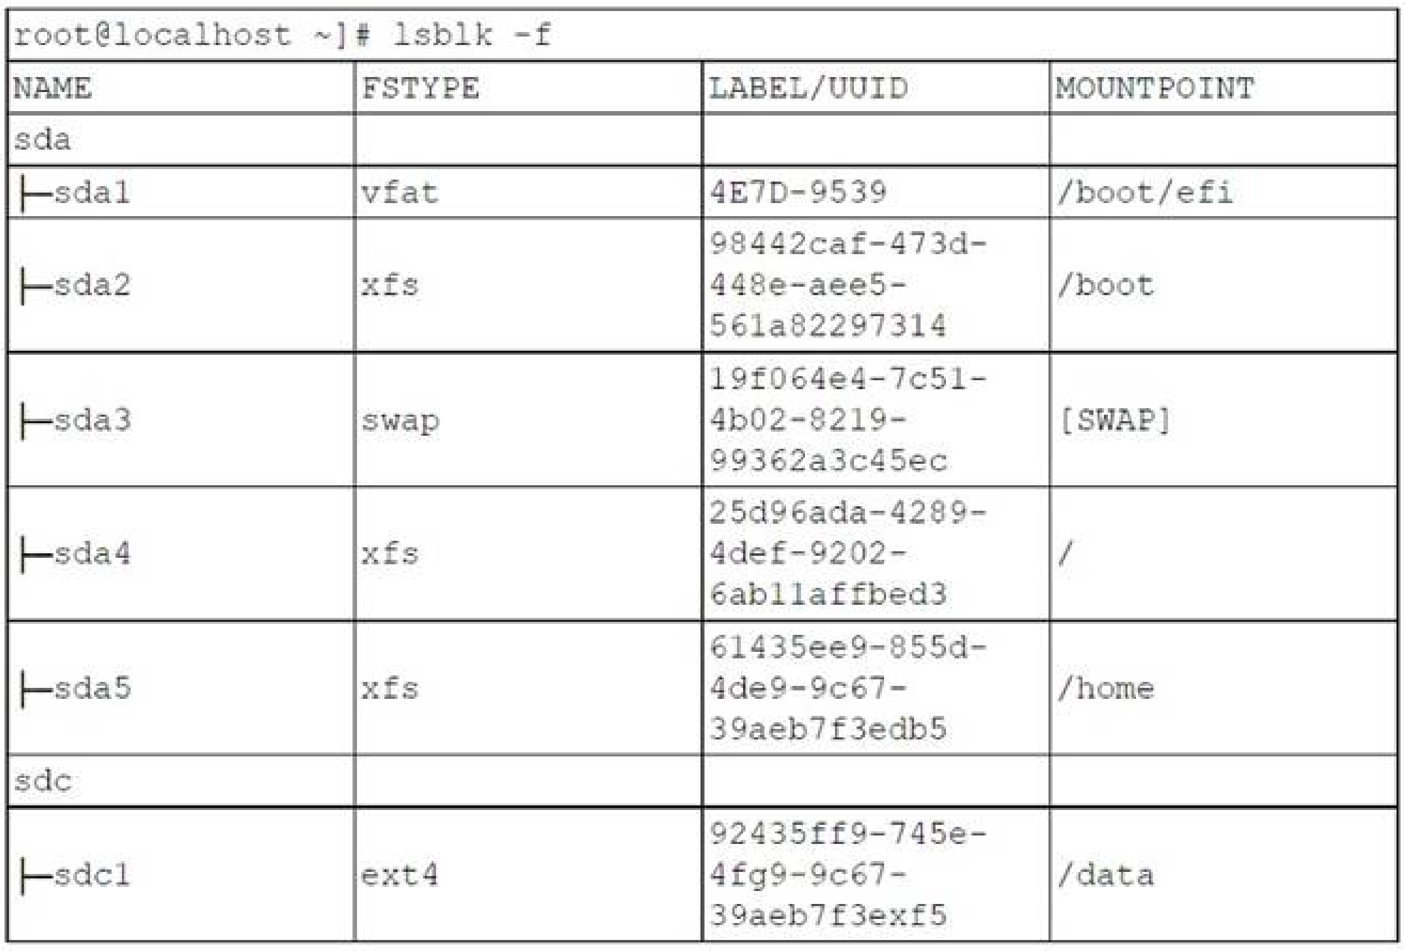

```
root@localhost ~]# lsblk -f
```

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

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

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

Answer: B

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

NEW QUESTION 229
A Linux administrator provisioned a new web server with custom administrative permissions for certain users. The administrator receives a report that user1 is unable to restart the Apache web service on this server. The administrator reviews the following output:
[root@server] # id user1
UID=1011 (user1) gid=1011 (USER1) groups=1011 (user1), 101 (www-data), 1120 (webadmin)
[root@server] # cat /etc/sudoers.d/custom.conf
user1 ALL=/usr/sbin/systemctl start httpd, /usr/sbin/systemctl stop httpd webadmin ALL=NOPASSWD: /etc/init.d.httpd restart, /sbin/service httpd restart, /usr/sbin/apache2ctl restart
#%wheel ALL=(ALL) NOPASSWD: ALL
Which of the following would most likely resolve the issue while maintaining a least privilege security model?

A. User1 should be added to the wheel group to manage the service.
B. User1 should have "NOPASSWD:" after the "ALL=" in the custo
C. conf.
D. The wheel line in the custo
E. conf file should be uncommented.
F. Webadmin should be listed as a group in the custo
G. conf file.

Answer: D

Explanation:
The custom.conf file grants sudo privileges to user1 and webadmin for managing the Apache web service, but it uses different commands for each of them. User1 is allowed to use systemctl to start and stop the httpd service, while webadmin is allowed to use init.d, service, or apache2ctl to restart the httpd service. However, the user1 is unable to restart the service, only start and stop it. To fix this, user1 should be able to use the same commands as webadmin, which can be achieved by listing webadmin as a group in the custom.conf file, using the syntax %groupname. This way, user1 will inherit the sudo privileges of the webadmin group, and be able to restart the Apache web service without compromising the least privilege security model.
References
? Sudo and Sudoers Configuration | Servers for Hackers, section “Groups”
? Chapter 12. Managing sudo access - Red Hat Customer Portal, section “12.1.
Configuring sudo access for users and groups”

NEW QUESTION 234
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual XK0-005 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the XK0-005 Product From:

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

Money Back Guarantee

XK0-005 Practice Exam Features:

* XK0-005 Questions and Answers Updated Frequently

* XK0-005 Practice Questions Verified by Expert Senior Certified Staff

* XK0-005 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* XK0-005 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year