CompTIA

Exam Questions XK0-005

CompTIA Linux+ Certification Exam

NEW QUESTION 1
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 2
A Linux administrator needs to obtain a list of all volumes that are part of a volume group. Which of the following commands should the administrator use to accomplish this task?

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

Answer: B

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

NEW QUESTION 3
The administrator comptia is not able to perform privileged functions on a newly deployed system. Given the following command outputs:

```
[root@newserver ~]# id comptia
uid=1000(comptia) gid=1000(comptia) groups=1000(comptia)

[root@newserver ~]# cat /etc/sudoers.d/admin
%admin ALL= (root) NOPASSWD: EXEC: /usr/bin/ps, /usr/bin/chmod, /usr/bin/yum, /usr/bin/cat, /usr/sbin/lvm,
/usr/sbin/pvs

[root@newserver ~]# grep comptia /etc/passwd
comptia:x:1000:1000:comptia:/home/comptia:/bin/bash

[root@newserver ~]# chage -l comptia
Last password change : never
Password expires : never
Password inactive : never
Account expires : never
Minimum number of days between password change : 0
Maximum number of days between password change : 99999
Number of days of warning before password expires : 7
```

Which of the following is the reason that the administrator is unable to perform the assigned duties?

A. The administrator needs a password reset.
B. The administrator is not a part of the correct group.
C. The administrator did not update the sudo database.
D. The administrator's credentials need to be more complex.

Answer: B

Explanation:
The reason that the administrator is unable to perform the assigned duties is because the administrator is not a part of the correct group. This is option B.
Based on the image that you sent, I can see that the user comptia has a user ID and a group ID of 1000, and belongs to only one group, which is also comptia. However, the sudoers file, which defines the permissions for users to run commands as root or other users, does not include the comptia group in any of the entries. Therefore, the user comptia cannot use sudo to perform privileged functions on the system.
The other options are incorrect because:
* A. The administrator needs a password reset.
This is not true, because the password aging information for the user comptia shows that the password was last changed on Oct 24, 2023, and it does not expire until Jan 22, 2024. There is no indication that the password is locked or expired.
* C. The administrator did not update the sudo database.
This is not necessary, because the sudo database is automatically updated whenever the sudoers file is modified. There is no separate command to update the sudo database.
* D. The administrator’s credentials need to be more complex.
This is not relevant, because the complexity of the credentials does not affect the ability to use sudo. The sudoers file does not specify any password policy for the users or groups that are allowed to use sudo.

NEW QUESTION 4
After starting an Apache web server, the administrator receives the following error:
Apr 23 localhost.localdomain httpd 4618] : (98) Address already in use: AH00072: make_sock: could not bind to address [: :]80
Which of the following commands should the administrator use to further trou-bleshoot this issue?

A. Ss
B. Ip
C. Dig
D. Nc

Answer: A

Explanation:
The ss command is used to display information about socket connections, such as the port number, state, and process ID. The error message indicates that the port 80 is already in use by another process, which prevents the Apache web server from binding to it. By using the ss command with the -l and -n options, the administrator can list all the listening sockets and their port numbers in numeric form, and identify which process is using the port 80. For example: ss -ln | grep :80. The ip, dig, and nc commands are not relevant for this issue, as they are used for different purposes, such as configuring network interfaces, querying DNS records, and testing network connectivity.

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

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

Answer: A

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

NEW QUESTION 6
A Linux administrator was asked to run a container with the httpd server inside. This container should be exposed at port 443 of a Linux host machine while it internally listens on port 8443. Which of the following commands will accomplish this task?

A. podman run -d -p 443:8443 httpd
B. podman run -d -p 8443:443 httpd
C. podman run –d -e 443:8443 httpd
D. podman exec -p 8443:443 httpd

Answer: A

Explanation:
The command that will accomplish the task of running a container with the httpd server inside and exposing it at port 443 of the Linux host machine while it internally listens on port 8443 is podman run -d -p 443:8443 httpd. This command uses the podman tool, which is a daemonless container engine that can run and manage containers on Linux systems. The -d option runs the container in detached mode, meaning that it runs in the background without blocking the terminal. The -p option maps a port on the host machine to a port inside the container, using the format host_port:container_port. In this case, port 443 on the host machine is mapped to port 8443 inside the container, allowing external access to the httpd server. The httpd argument specifies the name of the image to run as a container, which in this case is an image that contains the Apache HTTP Server software. The other options are not correct commands for accomplishing the task. Podman run -d -p 8443:443 httpd maps port 8443 on the host machine to port 443 inside the container, which does not match the requirement. Podman run –d -e 443:8443 httpd uses the -e option instead of the -p option, which sets an environment variable inside the container instead of mapping a port. Podman exec -p 8443:443 httpd uses the podman exec command instead of the podman run command, which executes a command inside an existing container instead of creating a new one. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks

NEW QUESTION 7
A systems technician is working on deploying several microservices to various RPM-based systems, some of which could run up to two hours. Which of the following commands will allow the technician to execute those services and continue deploying other microservices within the same terminal section?

A. gedit & disown
B. kill 9 %1
C. fg %1
D. bg %1 job name

Answer: D

Explanation:
The command that will allow the technician to execute the services and continue deploying other microservices within the same terminal session is bg %1 job name. This command will send the job with ID 1 and name job name to the background, where it will run without occupying the terminal. The other options are incorrect because:
? gedit & disown will launch a graphical text editor in the background and detach it from the terminal, but it will not execute any service.
? kill 9 %1 will terminate the job with ID 1 using a SIGKILL signal, which cannot be ignored or handled by the process.
? fg %1 will bring the job with ID 1 to the foreground, where it will occupy the terminal until it finishes or is stopped. References: CompTIA Linux+ Study Guide, Fourth Edition, page 181-182.

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

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

Answer: A

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

NEW QUESTION 9
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 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 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 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 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 10
A systems administrator is tasked with setting up key-based SSH authentication. In which of the following locations should the administrator place the public keys for the server?

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

Answer: C

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

NEW QUESTION 10
An administrator needs to make an application change via a script that must be run only in console mode. Which of the following best represents the sequence the administrator should execute to accomplish this task?

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

Answer: A

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

NEW QUESTION 14
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 17
A systems administrator frequently connects to a remote host via SSH and a non-standard port. The systems administrator would like to avoid passing the port parameter on the command line every time. Which of the following files can be used to set a different port
value for that host?

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

Answer: C

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

NEW QUESTION 21
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 26
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 28
Which of the following directories is the mount point in a UEFI system?

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

Answer: B

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

NEW QUESTION 31
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 34
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 35
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 36
The group owner of the / home/ test directory would like to preserve all group permissions on files created in the directory. Which of the following commands should the group owner execute?

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

Answer: A

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

octal number that represents the permissions in binary form.
* D. chown -hR test /home/test
This command will change the owner and group of the /home/test directory and all its contents recursively to test, but it will not preserve the original group permissions on files created in the directory. The chown command is used to change the owner and group of files and directories. The -hR option is used to affect symbolic links and operate on all files and directories recursively. The test /home/test arguments are used to specify the new owner and group and the target directory.
References:
? How to Set File Permissions Using chmod
? How to Use Chmod Command in Linux with Examples
? How to Use Chown Command in Linux with Examples
? [How to Use Chgrp Command in Linux with Examples]

NEW QUESTION 38
A DevOps engineer needs to download a Git repository from https://git.company.com/admin/project.git. Which of the following commands will achieve this goal?

A. git clone https://git.company.com/admin/project.git
B. git checkout https://git.company.com/admin/project.git
C. git pull https://git.company.com/admin/project.git
D. git branch https://git.company.com/admin/project.git

Answer: A

Explanation:
The command git clone https://git.company.com/admin/project.git will achieve the goal of downloading a Git repository from the given URL. The git command is a tool for managing version control systems. The clone option creates a copy of an existing repository. The URL specifies the location of the repository to clone, in this case https://git.company.com/admin/project.git. The command git clone https://git.company.com/admin/project.git will download the repository and create a directory named project in the current working directory. This is the correct command to use to accomplish the goal. The other options are incorrect because they either do not download the repository (git checkout, git pull, or git branch) or do not use the correct syntax (git checkout https://git.company.com/admin/project.git instead of git checkout -b project https://git.company.com/admin/project.git or git branch https://git.company.com/admin/project.git instead of git branch project https://git.company.com/admin/project.git). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 571.

NEW QUESTION 42
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 46
A Linux engineer set up two local DNS servers (10.10.10.10 and 10.10.10.20) and was testing email connectivity to the local mail server using the mail command on a local machine when the following error appeared:

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

The local machine DNS settings are:

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

Mail server: mail.example.com
```

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

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

Answer: B

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

NEW QUESTION 50
Users have been unable to save documents to /home/tmp/temp and have been receiving the following error:
Path not found
A junior technician checks the locations and sees that /home/tmp/tempa was accidentally created instead of /home/tmp/temp. Which of the following commands should the technician use to fix this issue?

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

Answer: B

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

NEW QUESTION 53
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 57
An administrator installed an application from source into /opt/operations1/ and has received numerous reports that users are not able to access the application without having to use the full path /opt/operations1/bin/*. Which of the following commands should be used to resolve this issue?

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

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

Answer: A

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

NEW QUESTION 62
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 65
A Linux administrator cloned an existing Linux server and built a new server from that clone. The administrator encountered the following error after booting the cloned server:

```
Device mismatch detected
```

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

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

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

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

Answer: A

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

NEW QUESTION 70
A Linux administrator needs to transfer a local file named accounts . pdf to a remote / tmp directory of a server with the IP address 10.10.10.80. Which of the following commands needs to be executed to transfer this file?

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

Answer: B

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

NEW QUESTION 73
A Linux engineer has been notified about the possible deletion of logs from the file
/opt/app/logs. The engineer needs to ensure the log file can only be written into without removing previous entries.

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

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

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

Answer: A

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

NEW QUESTION 78
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 79
A Linux administrator has set up a new DNS forwarder and is configuring all internal servers to use the new forwarder to look up external DNS requests. The

administrator needs to modify the firewall on the server for the DNS forwarder to allow the internal servers to communicate to it and make the changes persistent between server reboots. Which of the following commands should be run on the DNS forwarder server to accomplish this task?

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

Answer: D

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

NEW QUESTION 80
A Linux administrator needs to redirect all HTTP traffic temporarily to the new proxy server 192.0.2.25 on port 3128. Which of the following commands will accomplish this task?

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

Answer: D

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

NEW QUESTION 84
A Linux system is failing to boot with the following error:

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

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

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

Answer: BF

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

NEW QUESTION 89

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

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

Answer: B

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

NEW QUESTION 94
A Linux administrator needs to connect securely to a remote server in order to install application software. Which of the following commands would allow this connection?

A. scp "ABC-key.pem" root@10.0.0.1
B. sftp rooteiO.0.0.1
C. telnet 10.0.0.1 80
D. ssh -i "ABC-key.pem" root@10.0.0.1
E. sftp "ABC-key.pem" root@10.0.0.1

Answer: D

Explanation:
The command ssh -i "ABC-key.pem" root@10.0.0.1 would allow the administrator to connect securely to the remote server in order to install application software. The ssh command is a tool for establishing secure and encrypted connections between remote systems. The -i option specifies the identity file that contains the private key for key-based authentication. The "ABC-key.pem" is the name of the identity file that contains the private key. The root@10.0.0.1 is the username and the IP address of the remote server. The command ssh -i "ABC-key.pem" root@10.0.0.1 will connect to the remote server using the private key and allow the administrator to install application software. This is the correct command to use to connect securely to the remote server. The other options are incorrect because they either do not use key-based authentication (sftp root@10.0.0.1 or telnet 10.0.0.1 80) or do not use the correct syntax for the command (scp "ABC-key.pem" root@10.0.0.1 instead of scp -i "ABC-key.pem" root@10.0.0.1 or sftp "ABC-key.pem" root@10.0.0.1 instead of sftp -i "ABC-key.pem" root@10.0.0.1). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Implementing Basic Security, page 513.

NEW QUESTION 95
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 97
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 102
A Linux administrator is tasked with adding users to the system. However, the administrator wants to ensure the users' access will be disabled once the project is over. The expiration date should be 2021-09-30. Which of the following commands will accomplish this task?

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

Answer: A

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

NEW QUESTION 103
A systems administrator needs to remove a disk from a Linux server. The disk size is 500G, and it is the only one that size on that machine. Which of the following commands can the administrator use to find the corresponding device name?

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

Answer: D

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

NEW QUESTION 106
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 107
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 110
Employees in the finance department are having trouble accessing the file /opt/work/file. All IT employees can read and write the file. Systems administrator reviews the following output:

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

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

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

Answer: D

Explanation:
The command setfacl -m g:finance:rw file will permanently fix the access issue while limiting access to IT and finance department employees. The setfacl command is a tool for modifying the access control lists (ACLs) of files and directories on Linux systems. The ACLs are a mechanism that allows more fine-grained control over the permissions of files and directories than the traditional owner-group-others model. The - m option specifies the modification to the ACL. The g:finance:rw means that the group named finance will have read and write permissions on the file. The file is the name of the file to modify, in this case /opt/work/file. The command setfacl -m g:finance:rw file will add an entry to the ACL of the file that will grant read and write access to the finance group.
This will fix the access issue and allow the finance employees to access the file. The command will also preserve the existing permissions of the file, which means that the IT employees will still have read and write access to the file. This will limit the access to IT and finance department employees and prevent unauthorized access from other users.
This is the correct command to use to accomplish the task. The other options are incorrect because they either do not fix the access issue (chattr +i file or chown it:finance file) or do not limit the access to IT and finance department employees (chmod 666 file). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing File Permissions and Ownership, page 352.

NEW QUESTION 111
A Linux administrator has defined a systemd script docker-repository.mount to mount a volume for use by the Docker service. The administrator wants to ensure that Docker service does not start until the volume is mounted. Which of the following configurations needs to be added to the Docker service definition to best accomplish this task?

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

Answer: C

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

NEW QUESTION 114
An administrator would like to mirror the website files on the primary web server, www1, to the backup web server, www2. Which of the following commands should the administrator use to most efficiently accomplish this task?

A. [wwwl] rsync —a —e ssh /var/www/html/ user1@www2 : /var/www/html
B. [wwwl] scp —r /var/www/html user1@www2 : / var/www/html
C. [www2] cd /var/www/html; wget —m http: //wwwl/
D. [wwwl] cd /var/www/html && tar cvf —

Answer: A

Explanation:
To mirror the website files on the primary web server, www1, to the backup web server, www2, the administrator can use the command rsync -a -e ssh /var/www/html/ user1@www2:/var/www/html (A). This will synchronize all files and directories under /var/www/html/ on www1 to /var/www/html on www2 using ssh as the remote shell. The -a option will preserve all attributes and permissions of the files. The other commands will not mirror the website files, but either copy them once, download them from a web
server, or archive them. References:
? [CompTIA Linux+ Study Guide], Chapter 12: Troubleshooting Linux Systems, Section: Synchronizing Files with rsync
? [How to Use rsync Command in Linux]

NEW QUESTION 119
After listing the properties of a system account, a systems administrator wants to remove the expiration date of a user account. Which of the following commands will accomplish this task?

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

Answer: D

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

NEW QUESTION 124
A DevOps engineer needs to allow incoming traffic to ports in the range of 4000 to 5000 on a Linux server. Which of the following commands will enforce this rule?

A. iptables -f filter -I INPUT -p tcp --dport 4000:5000 -A ACCEPT
B. iptables -t filter -A INPUT -p tcp --dport 4000:5000 -j ACCEPT
C. iptables filter -A INPUT -p tcp --dport 4000:5000 -D ACCEPT
D. iptables filter -S INPUT -p tcp --dport 4000:5000 -A ACCEPT

Answer: B

Explanation:
The command iptables -t filter -A INPUT -p tcp --dport 4000:5000 -j ACCEPT will enforce the rule of allowing incoming traffic to ports in the range of 4000 to 5000 on a Linux server. The iptables command is a tool for managing firewall rules on Linux systems. The -t option specifies the table to operate on, in this case filter, which is the default table that contains the rules for filtering packets. The -A option appends a new rule to the end of a chain, in this case INPUT, which is the chain that processes the packets that are destined for the local system. The -p option specifies the protocol to match, in this case tcp, which is the transmission control protocol. The --dport option specifies the destination port or port range to match, in this case 4000:5000, which is the range of ports from 4000 to 5000. The -j option specifies the target to jump to if the rule matches, in this case ACCEPT, which is the target that allows the packet to pass through.
The command iptables -t filter -A INPUT -p tcp --dport 4000:5000 -j ACCEPT will add a new rule to the end of the INPUT chain that will accept the incoming TCP packets that have a destination port between 4000 and 5000. This command will enforce the rule and allow the traffic to the specified ports. This is the correct command to use to accomplish the task. The other options are incorrect because they either use the wrong options (-f instead of - t or -D instead of -A) or do not exist (iptables filter -A INPUT -p tcp --dport 4000:5000 -D ACCEPT or iptables filter -S INPUT -p tcp --dport 4000:5000 -A ACCEPT). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 543.

NEW QUESTION 127
A Linux administrator has been tasked with installing the most recent versions of packages on a RPM-based OS. Which of the following commands will accomplish this task?

A. apt-get upgrade
B. rpm -a
C. yum updateinfo
D. dnf update
E. yum check-update

Answer: D

Explanation:
The dnf update command will accomplish the task of installing the most recent versions of packages on a RPM-based OS. This command will check for available updates from the enabled repositories and apply them to the system. The apt-get upgrade command is used to install updates on a Debian-based OS, not a RPM-based OS. The rpm -a command is invalid, as -a is not a valid option for rpm. The yum updateinfo command will display information about available updates, but it will not install them. The yum check- update command will check for available updates, but it will not install them. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Packages and Software, page 559.

NEW QUESTION 131
A senior Linux administrator has created several scripts that will be used to install common system applications. These scripts are published to a reposito-ry to share with the systems team. A junior Linux administrator needs to re-trieve the scripts and make them available on a local workstation. Which of the following Git commands should the junior Linux administrator use to accom-plish this task?

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

Answer: C

Explanation:
To retrieve the scripts from a repository and make them available on a local workstation, the junior Linux administrator can use the command git clone ©. This will create a copy of the repository on the local machine, including all the scripts and history. The other commands will not clone the repository, but either fetch, checkout, or branch from an existing repository. References:
? [CompTIA Linux+ Study Guide], Chapter 10: Working with Git, Section: Cloning Repositories with Git
? [How to Clone a Git Repository]

NEW QUESTION 132
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 136
Which of the following technologies can be used as a central repository of Linux users and groups?

A. LDAP
B. MFA
C. SSO
D. PAM

Answer: A

Explanation:
LDAP stands for Lightweight Directory Access Protocol, which is a protocol for accessing and managing a central directory of users and groups. LDAP can be used as a central repository of Linux users and groups, allowing for centralized authentication and authorization across multiple Linux systems. MFA, SSO, and PAM are not technologies that can be used as a central repository of Linux users and groups. MFA stands for Multi- Factor Authentication, which is a method of verifying a user's identity using more than one factor, such as a password, a token, or a biometric. SSO stands for Single Sign-On, which is a feature that allows a user to log in once and access multiple applications or systems without having to re-enter credentials. PAM stands for Pluggable Authentication Modules, which is a framework that allows Linux to use different authentication methods, such as passwords, tokens, or biometrics. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing Users and Groups

NEW QUESTION 137
A Linux administrator is troubleshooting SSH connection issues from one of the workstations.
When users attempt to log in from the workstation to a server with the IP address 104.21.75.76, they receive the following message:

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

The administrator reviews the information below:

Workstation output 1:

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

Workstation output 2:

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

Server output 1:

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

Server output 2:

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

Server output 3:

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

Server output 4:

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

Which of the following is causing the connectivity issue?

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

Answer: C

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

NEW QUESTION 141
After installing some RPM packages, a systems administrator discovers the last package that was installed was not needed. Which of the following commands can be used to remove the package?

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

Answer: A

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

NEW QUESTION 142
At what point is the Internal Certificate Authority (ICA) created?

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

Answer: A

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

NEW QUESTION 144
The applications team is reporting issues when trying to access the web service hosted in a Linux system. The Linux systems administrator is reviewing the following outputs:
Output 1:
* httpd.service = The Apache HTTPD Server
Loaded: loaded (/usr/lib/systemd/system/httpd.service; disabled; vendor preset: disabled) Active: inactive (dead)
Docs: man:httpd(8) man:apachectl(8) Output 2:
16:51:16 up 28 min, 1 user, load average: 0.00, 0.00, 0.07
Which of the following statements best describe the root cause? (Select two).

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

Answer: CD

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

NEW QUESTION 149
A Linux administrator rebooted a server. Users then reported some of their files were missing. After doing some troubleshooting, the administrator found one of the filesystems was missing. The filesystem was not listed in /etc/f stab and might have been mounted manually by someone prior to reboot. Which of the following would prevent this issue from reoccurring in the future?

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

Answer: C

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

NEW QUESTION 152
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 156
An application developer received a file with the following content:
##This is a sample Image ## FROM ubuntu:18.04
MAINTAINER demohut@gtmail.com.hac COPY . /app
RUN make /app
CMD python /app/app.py RUN apt-get update
RUN apt-get install -y nginx CMD ["echo","Image created"]
The developer must use this information to create a test bed environment and identify the image (myimage) as the first version for testing a new application before moving it to production. Which of the following commands will accomplish this task?

A. docker build -t myimage:1.0 .
B. docker build -t myimage: .
C. docker build -t myimage-1.0 .
D. docker build -i myimage:1.0 .

Answer: A

Explanation:
The docker build command is used to build an image from a Dockerfile and a context1. The Dockerfile is a text file that contains the instructions for creating the image, and the context is a set of files that can be used in the image creation process1. The file that the developer received is an example of a Dockerfile.
The -t option is used to specify a name and an optional tag for the image1. The name and tag are separated by a colon (:), and the tag is usually used to indicate the version of the image2. For example, -t myimage:1.0 means that the image will be named myimage and tagged as 1.0.
The last argument of the docker build command is the path to the context, which can be a local directory or a URL1. The dot (.) means that the current working directory is the context2. Therefore, docker build -t myimage:1.0 . means that the image will be built from the Dockerfile and the files in the current working directory, and it will be named myimage and tagged as 1.0.

NEW QUESTION 160
A new disk was presented to a server as /dev/ sdd. The systems administrator needs to check if a partition table is on that disk. Which of the following commands can show this information?

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

Answer: B

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

NEW QUESTION 161
A Linux administrator is scheduling a system job that runs a script to check available disk space every hour. The Linux administrator does not want users to be able to start the job. Given the following:

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

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

[Install]
WantedBy=timers.target
```

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

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

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

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

Answer: C

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

NEW QUESTION 162
A cloud engineer wants to delete all unused networks that are not referenced by any container. Which of the following commands will achieve this goal?

A. docker network erase
B. docker network clear

C. docker network prune
D. docker network rm

Answer: C

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

NEW QUESTION 164
Which of the following will prevent non-root SSH access to a Linux server?

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

Answer: A

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

NEW QUESTION 167
A systems administrator is working on a security report from the Linux servers. Which of the following commands can the administrator use to display all the firewall rules applied to the Linux servers? (Select two).

A. ufw limit
B. iptables —F
C. systemct1 status firewalld
D. firewall—cmd ——1ist—a11
E. ufw status
F. iptables —A

Answer: DE

Explanation:
These commands can display all the firewall rules applied to the Linux servers, depending on which firewall service is being used.
? The firewall-cmd command is a utility for managing firewalld, which is a dynamic firewall service that supports zones and services. The --list-all option will show all the settings and rules for the default zone, or for a specific zone if specified. For example, firewall-cmd --list-all --zone=public will show the rules for the public zone1.
? The ufw command is a frontend for iptables, which is a low-level tool for manipulating netfilter, the Linux kernel's packet filtering framework. The status option will show the status of ufw and the active rules, or the numbered rules if verbose is specified. For example, ufw status verbose will show the numbered rules and other information2.
The other options are incorrect because:
* A. ufw limit
This command will limit the connection attempts to a service or port using iptables' recent module. It does not display any firewall rules2.
* B. iptables -F
This command will flush (delete) all the rules in the selected chain, or all chains if none is given. It does not display any firewall rules3.
* C. systemctl status firewalld
This command will show the status of the firewalld service, including whether it is active or not, but it does not show the firewall rules4.
* F. iptables -A
This command will append one or more rules to the end of the selected chain. It does not display any firewall rules3.

NEW QUESTION 169
A developer is trying to install an application remotely that requires a graphical interface for installation. The developer requested assistance to set up the necessary environment variables along with X11 forwarding in SSH. Which of the following environment variables must be set in remote shell in order to launch the graphical interface?

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

Answer: D

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

variable that stores the name of the current shell, but it is not related to X11 forwarding. References: How to enable or disable X11 forwarding in an SSH server; How to Configure X11 Forwarding Using SSH In Linux

NEW QUESTION 172
A Linux administrator needs to create a new user named user02. However, user02 must be in a different home directory, which is under /comptia/projects. Which of the following commands will accomplish this task?

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

Answer: A

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

NEW QUESTION 174
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 177
A systems engineer is adding a new 1GB XFS filesystem that should be temporarily mounted under /ops/app. Which of the following is the correct list of commands to achieve this goal?

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

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

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

D.

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

Answer: D

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

NEW QUESTION 181
A 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 184
Users have been unable to reach www.comptia.org from a Linux server. A systems administrator is troubleshooting the issue and does the following:

Output 1:
```
2: eth0: <BROADCAST,MULTICAST,UP, LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
   link/ether ac:11:22:33:44:cd brd ff:ff:ff:ff:ff:ff
   inet 192.168.168.10/24 brd 192.168.169.255 scope global dynamic noprefixroute eth0
      valid_lft 8097sec preferred_lft 8097sec
   inet fe80::4daf:8c7c:a6ff:2771/64 scope link noprefixroute
      valid_lft forever preferred_lft forever
```

Output 2:
```
nameserver 192.168.168.53
```

Output 3:
```
FING 192.168.168.53 (192.168.168.53) 56(84) bytes of data.
64 bytes from 192.168.168.53: icmp_seq=1 ttl=64 time=2.85 ms

--- 192.168.168.53 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 2.847/2.847/2.847/0.000 ms
```

Output 4:
```
192.168.168.0/24 dev eth0 proto kernel scope link src 192.168.168.10 metric 600
```

Output 5:
```
...
;; QUESTION SECTION:
;www.comptia.org. IN A

;; ANSWER SECTION:
. 0 CLASS4096 OPT 10 8 LgmNvk0AazU=

;; ADDITIONAL SECTION:
www.comptia.org. 3385 IN A 23.96.239.26
...
```

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

A. The name www.comptia.org does not point to a valid IP address.
B. The server 192.168.168.53 is unreachable.
C. No default route is set on the server.
D. The network interface eth0 is disconnected.

Answer: B

Explanation:
The issue is caused by the server 192.168.168.53 being unreachable. This server is the DNS server configured in the /etc/resolv.conf file, which is used to resolve domain names to IP addresses. The ping command shows that the server cannot be reached, and the nslookup command shows that the name www.comptia.org cannot be resolved using this server. The other options are incorrect because:

? The name www.comptia.org does point to a valid IP address, as shown by the nslookup command using another DNS server (8.8.8.8).
? The default route is set on the server, as shown by the ip route command, which shows a default gateway of 192.168.168.1.
? The network interface eth0 is connected, as shown by the ip link command, which shows a state of UP for eth0. References: CompTIA Linux+ Study Guide, Fourth Edition, page 457-458, 461-462.

NEW QUESTION 188
A junior administrator updated the PostgreSQL service unit file per the data-base administrator's recommendation. The service has been restarted, but changes have not been applied. Which of the following should the administrator run for the changes to take effect?

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

Answer: B

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

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

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

Answer: C

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

NEW QUESTION 197
Which of the following commands will display the operating system?

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

Answer: C

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

NEW QUESTION 202
A systems administrator receives reports that several virtual machines in a host are responding slower than expected. Upon further investigation, the administrator obtains the following output from one of the affected systems:

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

Which of the following best explains the reported issue?

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

Answer: D

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

NEW QUESTION 204
A systems administrator is investigating an issue in which one of the servers is not booting up properly. The journalctl entries show the following:

```
Sep 16 20:30:43 server kernel: acpi PNP0A03:00: _OSC failed (AE_NOT_FOUND);
-- Subject: Unit dev-mapper-centos\x2dapp.device has failed
Sep 16 20:32:15 server systemd[1]: Dependency failed for /opt/app
-- Subject: Unit opt-app.mount has failed
-- Unit opt-app.mount has failed
Sep 16 20:32:15 server systemd[1]: Dependency failed for Local File Systems.
-- Subject: Unit local-fs.target has failed
-- Unit local-fs.target has failed.
Sep 16 20:32:15 server systemd[1]: Dependency failed for Relabel all filesystem, if necessary.
-- Subject: Unit rhel-autorelabel.service has failed
-- Unit rhel-autorelabel.service has failed.
```

Which of the following will allow the administrator to boot the Linux system to normal mode quickly?

A. Comment out the /opt/app filesystem in /etc/fstab and reboot.
B. Reformat the /opt/app filesystem and reboot.
C. Perform filesystem checks on local filesystems and reboot.
D. Trigger a filesystem relabel and reboot.

Answer: A

Explanation:
The fastest way to boot the Linux system to normal mode is to comment out the /opt/app filesystem in /etc/fstab and reboot. This will prevent the system from trying to mount the /opt/app filesystem at boot time, which causes an error because the filesystem does not exist or is corrupted. Commenting out a line in /etc/fstab can be done by adding a # symbol at the beginning of the line. Rebooting the system will apply the changes and allow the system to boot normally. Reformatting the /opt/app filesystem will not help to boot the system, as it will erase any data on the filesystem and require manual intervention to create a new filesystem. Performing filesystem checks on local filesystems will not help to boot the system, as it will not fix the missing or corrupted /opt/app filesystem. Triggering a filesystem relabel will not help to boot the system, as it will only change the security context of files and directories according to SELinux policy. References: CompTIA Linux+ (XK0- 005) Certification Study Guide, Chapter 14: Managing Disk Storage, page 456.

NEW QUESTION 206
A Linux administrator created the directory /project/access2all. By creating this directory, the administrator is trying to avoid the deletion or modification of files from non-owners. Which of the following will accomplish this goal?

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

Answer: A

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

NEW QUESTION 208
A junior Linux administrator is tasked with installing an application. The installation guide states the application should only be installed in a run level 5 environment.

```
$ systemctl get-default
getty.target
```

Which of the following commands would ensure the server is set to runlevel 5?

A. systemct1 isolate multi-user.target
B. systemct1 isolate graphical.target
C. systemct1 isolate network.target
D. systemct1 isolate basic.target

Answer: B

Explanation:
The command that would ensure the server is set to runlevel 5 is systemct1 isolate graphical.target. This command will change the current target (or runlevel) of systemd to graphical.target, which is equivalent to runlevel 5 in SysV init systems. Graphical.target means that the system will start with a graphical user interface (GUI) and all services required for it.
The other options are not correct commands for setting the server to runlevel 5. The systemct1 isolate multi-user.target command will change the current target to multi- user.target, which is equivalent to runlevel 3 in SysV init systems. Multi-user.target means that the system will start with multiple user logins and networking, but without a GUI. The systemct1 isolate network.target command will change the current target to network.target, which is not a real runlevel but a synchronization point for network-related services. Network.target means that network functionality should be available, but does not specify whether it should be started before or after it. The systemct1 isolate basic.target command will change the current target to basic.target, which is also not a real runlevel but a synchronization point for basic system services. Basic.target means that all essential services should be started, but does not specify whether it should be started before or after it. References: systemd System and Service Manager; systemd.special(7) - Linux manual page

NEW QUESTION 209
A Linux administrator generated a list of users who have root-level command-line access to the Linux server to meet an audit requirement. The administrator analyzes the following /etc/passwd and /etc/sudoers files:
$ cat /etc/passwd
root:x: 0:0: :/home/root: /bin/bash lee: x: 500: 500: :/home/lee:/bin/tcsh
mallory:x: 501:501: :/root:/bin/bash
eve:x: 502: 502: /home/eve:/bin/nologin carl:x:0:503: :/home/carl:/bin/sh
bob:x: 504: 504: : /home/bob:/bin/ksh
alice:x: 505:505: :/home/alice:/bin/rsh
$ cat /etc/sudoers
Cmnd_Alias SHELLS = /bin/tcsh, /bin/sh, /bin/bash Cmnd_Alias SYSADMIN = /usr/sbin/tcpdump
ALL = (ALL) ALL
ALL = NOPASSWD: SYSADMIN
Which of the following users, in addition to the root user, should be listed in the audit report as having root-level command-line access? (Select two).

A. Carl
B. Lee
C. Mallory
D. Eve
E. Bob
F. Alice

Answer: AC

Explanation:
The users who have root-level command-line access are those who have either the same user ID (UID) as root, which is 0, or the ability to run commands as root using sudo. Based on the /etc/passwd and /etc/sudoers files, the users who meet these criteria are:
? Carl: Carl has the same UID as root, which is 0, as shown in the /etc/passwd file.
This means that Carl can log in as root and execute any command with root privileges1
? Mallory: Mallory has the ability to run commands as root using sudo, as shown in the /etc/sudoers file. The line ALL = (ALL) ALL means that any user can run any command as any other user, including root, by using sudo. Mallory can also use the root shell /bin/bash as her login shell, as shown in the /etc/passwd file2
Therefore, the correct answer is A and C. Lee, Eve, Bob, and Alice do not have root-level command-line access because they have different UIDs from root and they cannot use sudo to run commands as root. Lee can only use sudo to run the commands listed in the Cmnd_Alias SHELLS, which are /bin/tcsh, /bin/sh, and /bin/bash. Eve cannot log in at all because her login shell is /bin/nologin. Bob and Alice can only use sudo to run the command /usr/sbin/tcpdump without a password, as specified by the Cmnd_Alias SYSADMIN and the line ALL = NOPASSWD: SYSADMIN2

NEW QUESTION 212
A Linux administrator is troubleshooting the root cause of a high CPU load and average.

```
$ uptime
07:30:43 up 20 days, 3 min, 1 user, load average: 2.98, 3.62, 5.21

$ top
PID  USER PR  NI  VIRT RES   SHR  S %CPU %MEM TIME+  COMMAND
6295 user1 30 -10 5465 56465 8254 R 86.5 1.5 7:35.25 app1

$ ps -ef | grep user1
user1 6295 1 7:42:19 tty/1   06:48:29 /usr/local/bin/app1
```

Which of the following commands will permanently resolve the issue?

A. renice -n -20 6295
B. pstree -p 6295
C. iostat -cy 1 5
D. kill -9 6295

Answer: D

Explanation:
The command that will permanently resolve the issue of high CPU load and average is kill -9 6295. This command will send a SIGKILL signal to the process with the PID 6295, which is the process that is consuming 99.7% of the CPU according to the top output. The SIGKILL signal will terminate the process immediately and free up the CPU resources. The kill command is used to send signals to processes by PID or name.
The other options are not correct commands for resolving this issue. The renice -n -20 6295 command will change the priority (niceness) of the process with PID 6295 to -20, which is the highest priority possible. This will make the process more CPU-intensive, not less. The renice command is used to change the priority of

running processes. The pstree - p 6295 command will show a tree of processes with PID 6295 as the root. This will not affect the CPU load or average, but only display information. The pstree command is used to display a tree of processes. The iostat -cy 1 5 command will show CPU and disk I/O statistics for 5 iterations with an interval of 1 second. This will also not affect the CPU load or average, but only display information. The iostat command is used to report CPU and I/O statistics. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Troubleshooting Linux Systems; kill(1) - Linux manual page; renice(1) - Linux manual page; pstree(1) - Linux manual page; iostat(1) - Linux manual page

NEW QUESTION 217
A Linux user reported the following error after trying to connect to the system remotely: ssh: connect to host 10.0.1.10 port 22: Resource temporarily unavailable
The Linux systems administrator executed the following commands in the Linux system while trying to diagnose this issue:

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

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

Which of the following commands will resolve this issue?

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

Answer: C

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

NEW QUESTION 219
Which of the following data structures is written in JSON?
A)

```
---
name: user1
position: DevOps
floor: 3
```

B)

```
<table>
<tbody><tr>
<td>user1</td>
<td>DevOps</td>
<td>3</td>
</tr>
</tbody></table>
```

C)

```
<root>
   <floor>3</floor>
   <name>user1</name>
   <position>DevOps</position>
</root>
```

D)

```
{
   "name": "user1",
   "job": "DevOps",
   "floor": 3
}
```

A. Option A
B. Option B
C. Option C
D. Option D

Answer: C

Explanation:
Option C is the only data structure that is written in JSON format. JSON stands for JavaScript Object Notation, and it is a lightweight and human-readable data interchange format. JSON uses curly braces to enclose objects, which consist of key-value pairs separated by commas. JSON uses square brackets to enclose arrays, which consist of values separated by commas. JSON supports six data types: strings, numbers, booleans, null, objects, and arrays. Option C follows these rules and syntax of JSON, while the other options do not. Option A is written in XML format, which uses tags to enclose elements and attributes. Option B is written in YAML format, which uses indentation and colons to define key-value pairs. Option D is written in INI format, which uses sections and equal signs to define key-value pairs. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 21: Automating Tasks with Ansible, page 591.

NEW QUESTION 222
A developer reported an incident involving the application configuration file /etc/httpd/conf/httpd.conf that is missing from the server. Which of the following identifies the RPM package that installed the configuration file?

A. rpm -qf /etc/httpd/conf/httpd.conf
B. rpm -ql /etc/httpd/conf/httpd.conf
C. rpm —query /etc/httpd/conf/httpd.conf
D. rpm -q /etc/httpd/conf/httpd.conf

Answer: A

Explanation:
The rpm -qf /etc/httpd/conf/httpd.conf command will identify the RPM package that installed the configuration file. This command will query the database of installed packages and display the name of the package that owns the specified file. The rpm -ql /etc/httpd/conf/httpd.conf command is invalid, as -ql is not a valid option for rpm. The rpm --query /etc/httpd/conf/httpd.conf command is incorrect, as --query requires a package name, not a file name. The rpm -q /etc/httpd/conf/httpd.conf command is incorrect,
as -q requires a package name, not a file name. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Packages and Software, page 560.

NEW QUESTION 226
An administrator accidentally installed the httpd RPM package along with several dependencies. Which of the following options is the best way for the administrator to revert the package
installation?

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

Answer: D

Explanation:

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

NEW QUESTION 228
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

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