XK0-005 Dumps

CompTIA Linux+ Certification Exam

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

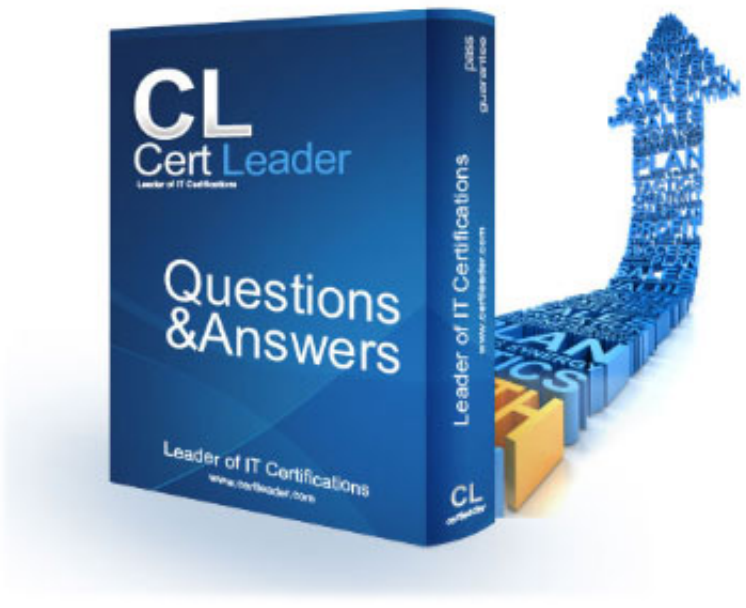

NEW QUESTION 1
A Linux administrator intends to start using KVM on a Linux server. Which of the following commands will allow the administrator to load the KVM module as well as any related dependencies?

A. modprobe kvm
B. insmod kvm
C. depmod kvm
D. hotplug kvm

Answer: A

Explanation:
This command will load the KVM module as well as any related dependencies, such as kvm-intel or kvm-amd, depending on the processor type. The modprobe command is a Linux utility that reads the /etc/modules.conf file and adds or removes modules from the kernel. It also resolves any dependencies between modules, so that they are loaded in the correct order.
The other options are incorrect because:
* B. insmod kvm
This command will only load the KVM module, but not any related dependencies. The insmod command is a low-level Linux utility that inserts a single module into the kernel. It does not resolve any dependencies between modules, so they have to be loaded manually.
* C. depmod kvm
This command will not load the KVM module at all, but only create a list of module dependencies for modprobe to use. The depmod command is a Linux utility that scans the installed modules and generates a file called modules.dep that contains dependency information for each module.
* D. hotplug kvm
This command is invalid and does not exist. The hotplug mechanism is a feature of the Linux kernel that allows devices to be added or removed while the system is running. It does not have anything to do with loading modules.

NEW QUESTION 2
A Linux administrator needs to 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
A systems administrator received a notification that a system is performing slowly. When running the top command, the systems administrator can see the following values:

```
%Cpu(s): 2.7 us, 1.9 sy, 0.0 ni, 0.4 id, 95 wa, 0.0 hi, 0.0 si 0.0 st
```

Which of the following commands will the administrator most likely run NEXT?

A. vmstat
B. strace
C. htop
D. lsof

Answer: A

Explanation:
The command vmstat will most likely be run next by the administrator to troubleshoot the system performance. The vmstat command is a tool for reporting virtual memory statistics on Linux systems. The command shows information about processes, memory, paging, block IO, interrupts, and CPU activity. The command can help the administrator identify the source of the performance issue, such as high CPU usage, low free memory, excessive swapping, or disk IO bottlenecks. The command can also be used with an interval and a count to display the statistics repeatedly over time and observe the changes. The command vmstat will provide useful information for diagnosing the system performance and finding the root cause of the issue. This is the most likely command to run next after the top command. The other options are incorrect because they either do not show the virtual memory statistics (strace or lsof) or do not provide more information than the top command (htop). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, page 425.

NEW QUESTION 4
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 5
A user reported issues when trying to log in to a Linux server. The following outputs were received:
Given the outputs above. which of the following is the reason the user is una-ble to log in to the server?

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

Answer: D

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

NEW QUESTION 6
A junior developer is unable to access an application server and receives the following output:

```
[root@server1 ~]# ssh dev2@172.16.25.126
dev2@172.16.25.126's password:
Permission denied, please try again.
dev2@172.16.25.126's password:
Permission denied, please try again.
dev2@172.16.25.126's password:
Account locked due to 4 failed logins
Account locked due to 5 failed logins
Last login: Mon Apr 22 21:21:06 2021 from 172.16.16.52
```

The systems administrator investigates the issue and receives the following output:

```
[root@server1 ~]# pam_tally2 --user=dev2
Login Failures Latest failure From
dev2 5 04/22/21 21:22:37 172.16.16.52
```

Which of the following commands will help unlock the account?

A. Pam_tally2 --user=dev2 —-quiet
B. pam_ tally2 --user=dev2
C. pam_tally2 -–user+dev2 —-quiet
D. pam_tally2 --user=dev2 —-reset

Answer: D

Explanation:
To unlock an account that has been locked due to login failures, the administrator can use the command pam_tally2 --user=dev2 --reset (D). This will reset the failure counter for the user “dev2” and allow the user to log in again. The other commands will not unlock the account, but either display or increase the failure count. References:
? [CompTIA Linux+ Study Guide], Chapter 4: Managing Users and Groups, Section: Locking Accounts with pam_tally2
? [How to Lock and Unlock User Account in Linux]

NEW QUESTION 7
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 8
A Linux administrator is troubleshooting a memory-related issue. Based on the output of the commands:

```
$ vmstat -s --unit M

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

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

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

Which of the following commands would address the issue?

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

Answer: B

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

NEW QUESTION 9
A cloud engineer is installing packages during VM provisioning. Which of the following should the engineer use to accomplish this task?

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

Answer: A

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

NEW QUESTION 10
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 10
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 11
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 13
Users are reporting that writes on a system configured with SSD drives have been taking longer than expected, but reads do not seem to be affected. A Linux systems administrator is investigating this issue and working on a solution. Which of the following should the administrator do to help solve the issue?

A. Run the corresponding command to trim the SSD drives.
B. Use fsck on the filesystem hosted on the SSD drives.
C. Migrate to high-density SSD drives for increased performance.
D. Reduce the amount of files on the SSD drives.

Answer: A

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

NEW QUESTION 18
The development team wants to prevent a file from being modified by all users in a Linux system, including the root account. Which of the following commands can be used to accomplish this objective?

A. chmod / app/conf/file
B. setenforce / app/ conf/ file
C. chattr +i /app/conf/file
D. chmod 0000 /app/conf/file

Answer: C

Explanation:
The chattr command is used to change file attributes on Linux systems that support extended attributes, such as ext2, ext3, ext4, btrfs, xfs, and others. File attributes are flags that modify the behavior of files and directories.
To prevent a file from being modified by all users in a Linux system, including the root account, the development team can use the chattr +i /app/conf/file command. This command will set the immutable attribute (+i) on the file /app/conf/file, which means that the file cannot be deleted, renamed, linked, appended, or written to by any user or process. To remove the immutable attribute, the development team can use the chattr -i /app/conf/file command. The statement C is correct.
The statements A, B, and D are incorrect because they do not prevent the file from being modified by all users. The chmod /app/conf/file command does not work because it requires an argument to specify the permissions to change. The setenforce /app/conf/file command does not work because it is used to change the SELinux mode, not file attributes. The chmod 0000 /app/conf/file command will remove all permissions from the file, but it can still be modified by the root account.
References: [How to Use chattr Command in Linux]

NEW QUESTION 21
A User on a Linux workstation needs to remotely start an application on a Linux server and then forward the graphical display of that application back to the Linux workstation. Which of the following would enable the user to perform this action?

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

Answer: A

Explanation:
The ssh -X option enables X11 forwarding, which allows the user to run graphical applications on the remote server and display them on the local workstation. The user needs to specify the username, the server address, and the application name after the ssh -X command. The remote server also needs to have

X11Forwarding enabled and xauth installed for this to work. References:
? The web search result 8 explains how to run a GUI application through SSH by configuring both the SSH client and server.
? The web search result 6 provides a detailed answer on how to forward X over SSH to run graphics applications remotely, with examples and troubleshooting tips.
? The CompTIA Linux+ Certification Exam Objectives mention that the candidate should be able to “use SSH for remote access and management” as part of the System Operation and Maintenance domain1.

NEW QUESTION 22
A Linux engineer needs to block an incoming connection from the IP address 2.2.2.2 to a secure shell server and ensure the originating IP address receives a response that a firewall is blocking the connection. Which of the following commands can be used to accomplish this task?

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

Answer: C

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

NEW QUESTION 24
A DevOps engineer wants to allow the same Kubernetes container configurations to be deployed in development, testing, and production environments. A key requirement is that the containers should be configured so that developers do not have to statically configure custom, environment-specific locations. Which of the following should the engineer use to meet this requirement?

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

Answer: D

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

NEW QUESTION 28
Which of the following can be used as a secure way to access a remote termi-nal?

A. TFTP
B. SSH
C. SCP
D. SFTP

Answer: B

Explanation:
SSH, or Secure Shell, is a protocol that allows you to access a remote terminal or virtual machine securely over an encrypted connection. You can use SSH to run commands, transfer files, or tunnel network traffic on a remote system. To use SSH, you need an SSH client program on your local system and an SSH server program on the remote system. You also need to authenticate yourself using a username and password or a public/private key pair. SSH is widely used by system administrators, developers, and engineers to remotely manage Linux servers and other devices.
The other options are not correct answers. TFTP, or Trivial File Transfer Protocol, is a simple protocol that allows you to transfer files between systems, but it does not provide any security or encryption features. SCP, or Secure Copy Protocol, is a protocol that uses SSH to securely copy files between systems, but it does not provide a remote terminal access. FTP, or File Transfer Protocol, is another protocol that allows you to transfer files between systems, but it also does not provide any security or encryption features.

NEW QUESTION 31
A developer has been unable to remove a particular data folder that a team no longer uses. The developer escalated the issue to the systems administrator. The following output was received:

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

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

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

Answer: C

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

NEW QUESTION 32
A systems administrator wants to delete app . conf from a Git repository. Which of the following commands will delete the file?

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

Answer: D

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

NEW QUESTION 35
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 36
A Linux system is getting an error indicating the root filesystem is full. Which of the following commands should be used by the systems administrator to resolve this issue? (Choose three.)

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

Answer: ACE

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

NEW QUESTION 38
A developer needs to launch an Nginx image container, name it Web001, and ex-pose port 8080 externally while mapping to port 80 inside the container. Which of the following commands will accomplish this task?

A. docker exec —it -p 8080: 80 ——name Web001 nginx
B. docker load -it -p 8080:80 ——name Web001 nginx
C. docker run -it -P 8080:80 ——name Web001 nginx
D. docker pull -it -p 8080:80 —name Web00l nginx

Answer: C

Explanation:
To launch an Nginx image container, name it Web001, and expose port 8080 externally while mapping to port 80 inside the container, the administrator can use the command docker run -it -p 8080:80 --name Web001 nginx ©. This will create and start a new container from the Nginx image, assign it a name of Web001, and map port 8080 on the host to port 80 on the container. The other commands are not valid or do not meet the requirements. References:
? [CompTIA Linux+ Study Guide], Chapter 11: Working with Containers, Section: Running Containers with Docker
? [How to Run Docker Containers]

NEW QUESTION 42
A Linux administrator is trying to remove the ACL from the file /home/user/data. txt but receives the following error message:

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

Given the following analysis:

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

-rw-rw-r--+ 1 user staff 2354 Sep 15 16:33 data.txt
-rw-rw-r--+ user staff unconfined_u:object_r:user_home_t:s0 data.txt

# file: data.txt
# owner: user
# group: staff
user::rw-
user:accounting:rw-
group::r-
mask::rw-
other::r-

Attributes:
-----a-----------
```

Which of the following is causing the error message?

A. The administrator is not using a highly privileged account.
B. The filesystem is mounted with the wrong options.
C. SELinux file context is denying the ACL changes.
D. File attributes are preventing file modification.

Answer: D

Explanation:
File attributes are preventing file modification, which is causing the error message. The output of lsattr /home/user/data.txt shows that the file has the immutable attribute (i) set, which means that the file cannot be changed, deleted, or renamed. The command setfacl -b /home/user/data.txt tries to remove the ACL from the file, but fails because of the immutable attribute. The administrator needs to remove the immutable attribute first by using the command chattr -i /home/user/data.txt and then try to remove the ACL again. The other options are incorrect because they are not supported by the outputs. The administrator is using a highly privileged account, as shown by the # prompt. The filesystem is mounted with the correct options, as shown by the output of mount | grep /home. SELinux file context is not denying the ACL changes, as shown by the output of ls - Z /home/user/data.txt. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, pages 357-358.

NEW QUESTION 46
A Linux administrator wants to set the SUID of a file named dev_team.text with 744 access rights. Which of the following commands will achieve this goal?

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

Answer: A

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

NEW QUESTION 51
A systems administrator is tasked with preventing logins from accounts other than root, while the file /etc/nologin exists. Which of the following PAM modules will accomplish this task?

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

Answer: D

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

NEW QUESTION 54
An administrator needs to make some changes in the IaC declaration templates. Which of the following commands would maintain version control?

A. git clone https://github.com/comptia/linux+-.git git push origin
B. git clone https://qithub.com/comptia/linux+-.git git fetch New-Branch

C. git clone https://github.com/comptia/linux+-.git git status
D. git clone https://github.com/comptia/linuxt+-.git git checkout -b <new-branch>

Answer: D

Explanation:
The command that will maintain version control while making some changes in the IaC declaration templates is git checkout -b <new-branch>. This command uses the git tool, which is a distributed version control system that tracks changes in source code and enables collaboration among developers. The checkout option switches to a different branch in the git repository, where a branch is a pointer to a specific commit in the history. The -b option creates a new branch with the given name, and switches to it. This way, the administrator can make changes in the new branch without affecting the main branch, and later merge them if needed.
The other options are not correct commands for maintaining version control while making some changes in the IaC declaration templates. The git clone https://github.com/comptia/linux±.git command will clone an existing repository from a remote URL to a local directory, but it will not create a new branch for making changes. The git push origin command will push the local changes to a remote repository named origin, but it will not create a new branch for making changes. The git fetch New-Branch command will fetch updates from a remote branch named New-Branch, but it will not create a new branch for making changes.
References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Source Code; Git - Basic Branching and Merging

NEW QUESTION 58
A Linux administrator found many containers in an exited state. Which of the following commands will allow the administrator to clean up the containers in an exited state?

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

Answer: B

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

NEW QUESTION 60
Using AD Query, the security gateway connections to the Active Directory Domain Controllers using what protocol?

A. Windows Management Instrumentation (WMI)
B. Hypertext Transfer Protocol Secure (HTTPS)
C. Lightweight Directory Access Protocol (LDAP)
D. Remote Desktop Protocol (RDP)

Answer: C

Explanation:
Using AD Query, the security gateway connects to the Active Directory Domain Controllers using Lightweight Directory Access Protocol (LDAP). LDAP is a protocol that provides access to directory services over a network. AD Query uses LDAP queries to retrieve information about users and groups from Active Directory Domain Controllers without installing any software on them. AD Query does not use Windows Management Instrumentation (WMI), Hypertext Transfer Protocol Secure (HTTPS), or Remote Desktop Protocol (RDP) to connect to Active Directory Domain Controllers. References: Check Point Certified Security Administrator (CCSA) R80.x Study Guide, Chapter 5: User Management and Authentication, page 69.

NEW QUESTION 62
An administrator is trying to diagnose a performance issue and is reviewing the following output:

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

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

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

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

Answer: B

Explanation:

The system has reached its maximum permitted throughput, therefore iowait is increasing. The output of iostat -x shows that the device sda has an average throughput of 44.01 MB/s, which is equal to the disk maximum throughput of 44 Mbps. The output also shows that the device sda has an average iowait of 99.99%, which means that the CPU is waiting for the disk to complete the I/O requests. This indicates that the disk is the bottleneck and the system is slow due to the high iowait. The other options are incorrect because they are not supported by the outputs. The system has not reached its maximum IOPS, as the device sda has an average IOPS of 563.50, which is lower than the disk maximum IOPS of 690. The system is not mostly idle, as the output of top shows that the CPU is 100% busy. The system does not have a partitioned disk, as the output of lsblk shows that the device sda has only one partition sda1. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Optimizing Linux Systems, pages 513-514.

NEW QUESTION 65
A systems administrator is tasked with mounting a USB drive on a system. The USB drive has a single partition, and it has been mapped by the system to the device /dev/sdb. Which of the following commands will mount the USB to /media/usb?

A. mount /dev/sdb1 /media/usb
B. mount /dev/sdb0 /media/usb
C. mount /dev/sdb /media/usb
D. mount -t usb /dev/sdb1 /media/usb

Answer: A

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

NEW QUESTION 67
DRAG DROP
A new drive was recently added to a Linux system. Using the environment and tokens provided, complete the following tasks:
• Create an appropriate device label.
• Format and create an ext4 file system on the new partition. The current working directory is /.

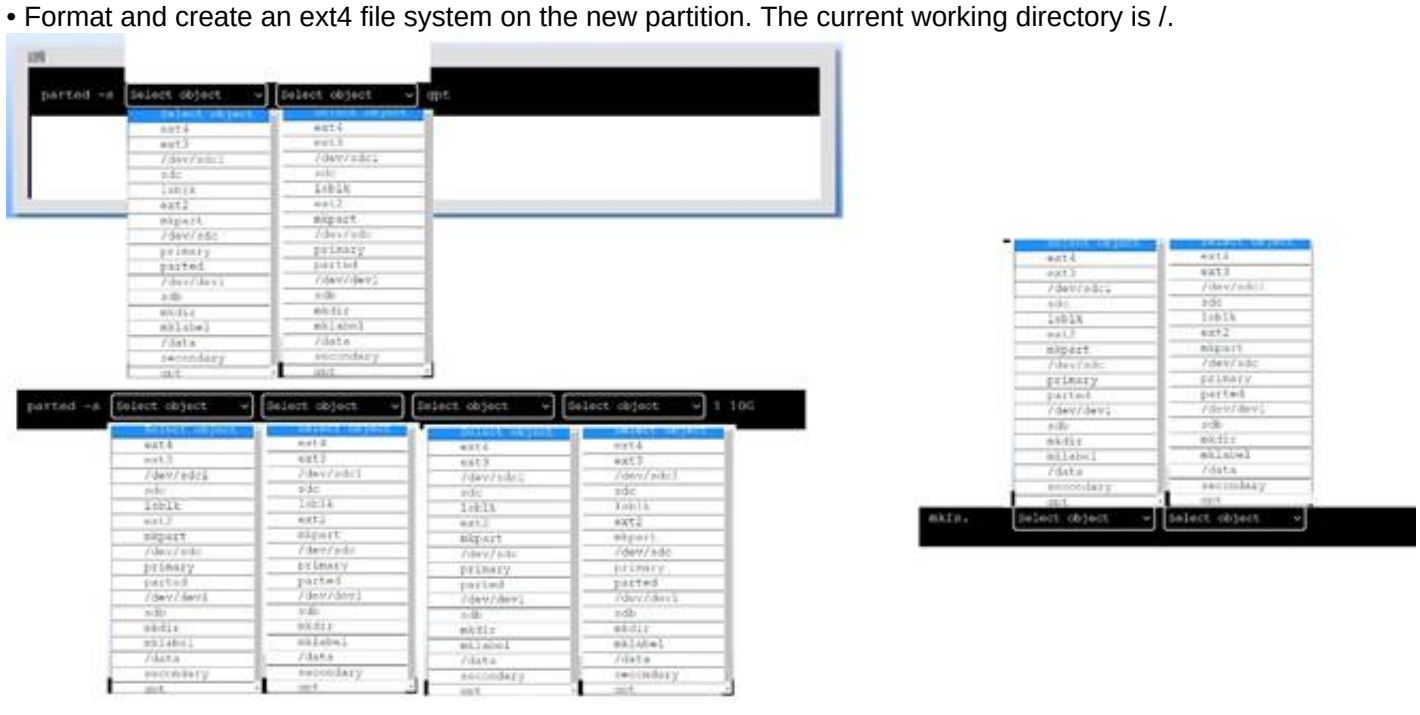

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 71
A Linux administrator needs to determine whether a hostname is in the DNS. Which of the following would supply the information that is needed?

A. nslookup
B. rsyn
C. netstat

D. host

Answer: A

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

NEW QUESTION 76
An administrator attempts to connect to a remote server by running the following command:
$ nmap 192.168.10.36
Starting Nmap 7.60 (https://nmap.org) at 2022-03-29 20:20 UTC Nmap scan report for www1 (192.168.10.36)
Host is up (0.000091s latency). Not shown: 979 closed ports PORT STATE SERVICE 21/tcp open ftp 22/tcp filtered ssh 631/tcp open ipp
Nmap done: 1 IP address (1 host up) scanned in 0.06 seconds
Which of the following can be said about the remote server?

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

Answer: A

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

NEW QUESTION 80
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 85
An administrator needs to increase the system priority of a process with PID 2274. Which of the following commands should the administrator use to accomplish this task?

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

Answer: A

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

NEW QUESTION 89
A Linux administrator needs to redirect all HTTP traffic temporarily to the new proxy server 192.0.2.25 on port 3128. Which of the following commands will accomplish this task?

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

Answer: D

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

NEW QUESTION 91
A 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 95
A Linux administrator is creating a primary partition on the replacement hard drive for an application server. Which of the following commands should the administrator issue to verify the device name of this partition?

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

Answer: C

Explanation:
The command sudo fdisk -l should be issued to verify the device name of the partition. The sudo command allows the administrator to run commands as the superuser or another user. The fdisk command is a tool for manipulating disk partitions on Linux systems. The -l option lists the partitions on all disks or a specific

disk. The command sudo fdisk -l will show the device names, sizes, types, and other information of the partitions on all disks. The administrator can identify the device name of the partition by looking at the output. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not list the partitions (sudo fdisk /dev/sda or sudo fdisk -h) or do not exist (sudo fdisk -s /dev/sda). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 317.

NEW QUESTION 100
A systems administrator received a request to change a user's credentials. Which of the following commands will grant the request?

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

Answer: A

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

NEW QUESTION 103
A Linux administrator reviews a set of log output files and needs to identify files that contain any occurrence of the word denied. All log files containing entries in uppercase or lowercase letters should be included in the list. Which of the following commands should the administrator use to accomplish this task?

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

Answer: D

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

NEW QUESTION 107
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 108
A Linux administrator recently downloaded a software package that is currently in a compressed file. Which of the following commands will extract the files?

A. unzip -v
B. bzip2 -z
C. gzip
D. funzip

Answer: C

Explanation:
The command gzip can extract files that are compressed with the gzip format, which has the extension .gz. This is the correct command to use for the software package. The other options are incorrect because they either compress files (bzip2 -z), unzip files that are compressed with the zip format (unzip -v or funzip), or have the wrong options (-v or -z instead of -d). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, page 353.

NEW QUESTION 109
An administrator accidentally deleted the /boot/vmlinuz file and must resolve the issue before the server is rebooted. Which of the following commands should the administrator use to identify the correct
version of this file?

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

Answer: A

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

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

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

Answer: D

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

NEW QUESTION 115
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 119
Users have reported that the interactive sessions were lost on a Linux server. A Linux administrator verifies the server was switched to rescue.target mode for maintenance. Which of the following commands will restore the server to its usual target?

A. telinit 0
B. systemct1 reboot
C. systemct1 get-default

D. systemct1 emergency

Answer: B

Explanation:
The systemct1 reboot command will restore the server to its usual target by rebooting it. This will cause the server to load the default target specified in /etc/systemd/system.conf or /etc/systemd/system/default.target files. The telinit 0 command would shut down the server, not restore it to its usual target. The systemct1 get-default command would display the default target, not change it. The systemct1 emergency command would switch the server to emergency.target mode, which is even more
restrictive than rescue.target mode. References: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 17: System Maintenance and Operation, page 516.

NEW QUESTION 121
A junior systems administrator has just generated public and private authentication keys for passwordless login. Which of the following files will be moved to the remote servers?

A. id_dsa.pem
B. id_rsa
C. id_ecdsa
D. id_rsa.pub

Answer: D

Explanation:
The file id_rsa.pub will be moved to the remote servers for passwordless login. The id_rsa.pub file is the public authentication key that is generated by the ssh-keygen command. The public key can be copied to the remote servers by using the ssh- copy-id command or manually. The remote servers will use the public key to authenticate the user who has the corresponding private key (id_rsa). This will allow the user to log in without entering a password. The other options are incorrect because they are either private keys (id_rsa, id_dsa.pem, or id_ecdsa) or non-existent files (id_dsa.pem or id_ecdsa). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 410.

NEW QUESTION 122
A systems administrator is tasked with changing the default shell of a system account in order to disable iterative logins. Which of the following is the best option for the administrator to use as the new shell?

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

Answer: A

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

NEW QUESTION 124
A systems administrator is installing various software packages using a pack-age manager. Which of the following commands would the administrator use on the Linux server to install the package?

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

Answer: D

NEW QUESTION 128
A Linux administrator needs to create a new cloud.cpio archive containing all the files from the current directory. Which of the following commands can help to accomplish this task?

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

Answer: C

Explanation:
The command ls | cpio -ov > cloud.cpio can help to create a new cloud.cpio archive containing all the files from the current directory. The ls command lists the files in the current directory and outputs them to the standard output. The | operator pipes the output to the next command. The cpio command is a tool for creating and extracting compressed archives. The -o option creates a new archive and the -v option shows the verbose output. The > operator redirects the output to the cloud.cpio file. This command will create a new cloud.cpio archive with all the files from the current directory. The other options are incorrect because they either

use the wrong options (-i instead of -o), the wrong arguments (cloud.epio instead of cloud.cpio), or the wrong syntax (< instead of > or missing |). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, page 351.

NEW QUESTION 132
A Linux administrator needs to expand a volume group using a new disk. Which of the following options presents the correct sequence of commands to accomplish the task?

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

Answer: D

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

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

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

Answer: A

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

NEW QUESTION 138
A Linux administrator is trying to start the database service on a Linux server but is not able to run it. The administrator executes a few commands and receives the following output:

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

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

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

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

A. systemctl unmask mariadb
B. journalctl —g mariadb
C. dnf reinstall mariadb
D. systemctl start mariadb
E. chkconfig mariadb on
F. service mariadb reload

Answer: AD

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

NEW QUESTION 142
An administrator created an initial Git repository and uploaded the first files. The administrator sees the following when listing the repository:

```
__init__.py      Initial Commit    Just now
main.py          Initial Commit    Just now
.DS_STORE        Initial Commit    Just now
setup.sh         Initial Commit    Just now
README.md        Initial Commit    Just now
```

The administrator notices the file . DS STORE should not be included and deletes it from the online repository. Which of the following should the administrator run from the root of the local repository before the next commit to ensure the file is not uploaded again in future commits?

A. rm -f .DS STORE && git push
B. git fetch && git checkout .DS STORE
C. rm -f .DS STORE && git rebase origin main
D. echo .DS STORE >> .gitignore

Answer: D

Explanation:
The correct answer is D. The administrator should run “echo .DS STORE >> .gitignore” from the root of the local repository before the next commit to ensure the file is not uploaded again in future commits.
This command will append the file name .DS STORE to the end of the .gitignore file, which is a special file that tells Git to ignore certain files or directories that should not be tracked or uploaded to the repository. By adding .DS STORE to the .gitignore file, the administrator will prevent Git from staging, committing, or pushing this file in the future.
The other options are incorrect because:
* A. rm -f .DS STORE && git push
This command will delete the file .DS STORE from the local repository and then push the changes to the remote repository. However, this does not prevent the file from being uploaded again in future commits, if it is recreated or copied to the local repository.
* B. git fetch && git checkout .DS STORE
This command will fetch the latest changes from the remote repository and then restore the file .DS STORE from the remote repository to the local repository. This is not what the administrator wants to do, as this will undo the deletion of the file from the online repository.
* C. rm -f .DS STORE && git rebase origin main
This command will delete the file .DS STORE from the local repository and then rebase the local branch onto the main branch of the remote repository. This will rewrite the commit history of the local branch and may cause conflicts or errors. This is not what the administrator wants to do, as this is a risky and unnecessary operation.

NEW QUESTION 145
A Linux administrator is configuring a two-node cluster and needs to be able to connect the nodes to each other using SSH keys from the root account. Which of the following commands will accomplish this task?

A. [root@nodea ssh —i ~/ . ssh/±d rsa root@nodeb
B. [root@nodea scp -i . ssh/id rsa root@nodeb
C. [root@nodea ssh—copy-id —i .ssh/id rsa root@nodeb
D. [root@nodea # ssh add -c ~/ . ssh/id rsa root@nodeb
E. [root@nodea # ssh add -c ~/. ssh/id rsa root@nodeb

Answer: C

Explanation:
The ssh-copy-id command is used to copy a public SSH key from a local machine to a remote server and add it to the authorized_keys file, which allows passwordless authentication between the machines. The administrator can use this command to copy the root user’s public key from nodea to nodeb, and vice versa, to enable SSH access between the nodes without entering a password every time. For example: [root@nodea ssh-copy-id -i ~/.ssh/id_rsa root@nodeb].
The ssh command is used to initiate an SSH connection to a remote server, but it does not copy any keys. The scp command is used to copy files securely between machines using SSH, but it does not add any keys to the authorized_keys file. The ssh-add command is used to add private keys to the SSH agent, which manages them for SSH authentication, but it does not copy any keys to a remote server.

NEW QUESTION 147
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 148
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 150
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 151
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 155
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 156
A Linux administrator needs to ensure that Java 7 and Java 8 are both locally available for developers to use when deploying containers. Currently only Java 8 is available. Which of the following commands should the administrator run to ensure both versions are available?

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

Answer: B

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

NEW QUESTION 160
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 164
A systems administrator is configuring a Linux system so the network traffic from the internal network 172.17.0.0/16 going out through the eth0 interface would appear as if it was sent directly from this interface. Which of the following commands will accomplish this task?

A. iptables -A POSTROUTING -s 172.17.0.0/16 -o eth0 -j MASQUERADE
B. firewalld -A OUTPUT -s 172.17.0.0/16 -o eth0 -j DIRECT
C. nmcli masq-traffic eth0 -s 172.17.0.0/16 -j MASQUERADE
D. ifconfig -- nat eth0 -s 172.17.0.0/16 -j DIRECT

Answer: A

Explanation:
This command will use the iptables tool to append a rule to the POSTROUTING chain of the nat table, which will match any packet with a source address of 172.17.0.0/16 and an output interface of eth0, and apply the MASQUERADE target to it. This means that the packet will have its source address changed to the address of the eth0 interface, effectively hiding the internal network behind a NAT12.
References: 1: Iptables NAT and Masquerade rules - what do they do? 2: Routing from docker containers using a different physical network interface and default gateway

NEW QUESTION 165
A Linux engineer receives reports that files created within a certain group are being modified by users who are not group members. The engineer wants to reconfigure the server so that only file owners and group members can modify new files by default. Which of the following commands would accomplish this task?

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

Answer: B

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

NEW QUESTION 166
The security team has identified a web service that is running with elevated privileges A Linux administrator is working to change the systemd service file to meet security compliance standards. Given the following output:

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

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

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

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

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

Answer: C

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

NEW QUESTION 171
A systems administrator wants to list all local accounts in which the UID is greater than 500. Which of the following commands will give the correct output?

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

Answer: C

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

NEW QUESTION 174
A Linux administrator modified the SSH configuration file. Which of the following commands should be used to apply the configuration changes?

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

Answer: C

Explanation:
The systemct1 reload sshd command can be used to apply the configuration changes of the SSH server daemon without restarting it. This is useful to avoid interrupting existing connections. The systemct1 stop sshd command would stop the SSH server daemon, not apply the changes. The systemct1 mask sshd

command would prevent the SSH server daemon from being started, not apply the changes. The systemct1 start sshd command would start the SSH server daemon if it is not running, but it would not apply the changes if it is already running. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Secure Shell (SSH), page 415.

NEW QUESTION 178
A cloud engineer needs to remove all dangling images and delete all the images that do not have an associated container. Which of the following commands will help to accomplish this task?

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

Answer: A

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

NEW QUESTION 180
A systems administrator is troubleshooting connectivity issues and trying to find out why a Linux server is not able to reach other servers on the same subnet it is connected to. When listing link parameters, the following is presented:

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

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

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

Answer: D

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

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

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

Answer: C

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

NEW QUESTION 185
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 186
A Linux administrator booted up the server and was presented with a non-GUI terminal. The administrator ran the command systemct1 isolate graphical.target and rebooted the system by running systemct1 reboot, which fixed the issue. However, the next day the administrator was presented again with a non-GUI terminal. Which of the following is the issue?

A. The administrator did not reboot the server properly.
B. The administrator did not set the default target to basic.target.
C. The administrator did not set the default target to graphical.target.
D. The administrator did not shut down the server properly.

Answer: C

Explanation:
The issue is that the administrator did not set the default target to graphical.target. A target is a unit of systemd that groups together other units by a common purpose or state. The graphical.target is a target that starts the graphical user interface (GUI) along with other services. The administrator used the command systemct1 isolate graphical.target to switch to this target temporarily, but this does not change the default target that is activated at boot time. To make this change permanent, the administrator should have used the command systemct1 set-default graphical.target, which creates a symbolic link from /etc/systemd/system/default.target to /usr/lib/systemd/system/graphical.target.
The other options are not correct explanations for the issue. The administrator did reboot the server properly by using systemct1 reboot, which shuts down and restarts the system cleanly. The administrator did not need to set the default target to basic.target, which is a minimal target that only starts essential services. The administrator did not shut down the server improperly, which could have caused file system corruption or data loss, but not affect the default target. References: systemct1(1) - Linux manual page; How to Change Runlevels (targets) in SystemD

NEW QUESTION 190
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 194
After installing a new version of a package, a systems administrator notices a new version of the corresponding, service file was Installed In order to use the new version of the, service file, which of the following commands must be Issued FIRST?

A. systemct1 status
B. systemct1 stop
C. systemct1 reinstall
D. systemct1 daemon-reload

Answer: D

Explanation:
After installing a new version of a package that includes a new version of the corresponding service file, the systemct1 daemon-reload command must be issued first in order to use the new version of the service file. This command will reload the systemd manager configuration and read all unit files that have changed on disk. This will ensure that systemd recognizes the new service file and applies its settings correctly. The systemct1 status command will display information about a service unit, but it will not reload the configuration. The systemct1 stop command will stop a service unit, but it will not reload the configuration. The systemct1 reinstall command does not exist. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: System Maintenance and Operation, page 518.

NEW QUESTION 199
A new file was added to a main Git repository. An administrator wants to synchronize a local copy with the contents of the main repository. Which of the following commands should the administrator use for this task?

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

Answer: B

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

NEW QUESTION 200
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 203
A user is unable to log on to a Linux workstation. The systems administrator executes the following command:
cat /etc/shadow | grep user1
The command results in the following output:
user1 :! 6QERgAsdvojadv4asdvaarC/9dj34GdafGVaregmkdsfa:18875:0:99999:7 ::: Which of the following should the systems administrator execute to fix the issue?

A. chown -R userl:user1 /home/user1
B. sed -i '/ ::: / :: /g' /etc/shadow
C. chgrp user1:user1 /home/user1
D. passwd -u user1

Answer: D

Explanation:
The output shows that the user1 account has a locked password, indicated by the exclamation point (!) in the second field of the /etc/shadow file1. To unlock the password and allow the user to log in, the systems administrator should use the passwd command with the -u (unlock) option2.
References: 1: Understanding the /etc/shadow File 2: How To Use The Passwd Command In Linux

NEW QUESTION 205
After connecting to a remote host via SSH, an administrator attempts to run an application but receives the following error:
[user@workstation ~]$ ssh admin@srv1 Last login: Tue Mar 29 18:03:34 2022
[admin@srvl ~] $ /usr/local/bin/config_manager Error: cannot open display:
[admin@srv1 ~] $
Which of the following should the administrator do to resolve this error?

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

Answer: A

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

NEW QUESTION 210
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 212
Which of the following tools is BEST suited to orchestrate a large number of containers across many different servers?

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

Answer: A

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

NEW QUESTION 216
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 220
Which of the following is a function of a bootloader?

A. It initializes all the devices that are required to load the OS.
B. It mounts the root filesystem that is required to load the OS.
C. It helps to load the different kernels to initiate the OS startup process.
D. It triggers the start of all the system services.

Answer: C

Explanation:
A function of a bootloader is to help load the different kernels to initiate the OS startup process. A bootloader is a program that runs when the system is powered on and prepares the system for booting the OS. A bootloader can load different kernels, which are the core components of the OS, and pass the control to the selected kernel. A bootloader can also provide a menu for the user to choose which kernel or OS to boot. This is a correct function of a bootloader. The other options are incorrect because they are
either functions of the kernel (initialize devices or mount root filesystem) or functions of the init system (trigger the start of system services). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing the Linux Boot Process, page 265.

NEW QUESTION 225
A systems administrator made some unapproved changes prior to leaving the company. The newly hired administrator has been tasked with revealing the system to a compliant state. Which of the following commands will list and remove the correspondent packages?

A. dnf list and dnf remove last
B. dnf remove and dnf check
C. dnf info and dnf upgrade
D. dnf history and dnf history undo last

Answer: D

Explanation:
The commands that will list and remove the corresponding packages are dnf history and dnf history undo last. The dnf history command will display a list of all transactions performed by dnf, such as installing, updating, or removing packages. Each transaction has a unique ID, a date and time, an action, and a number of altered packages. The dnf history undo last command will undo the last transaction performed by dnf, meaning that it will reverse all package changes made by that transaction. For example, if the last transaction installed some packages, dnf history undo last will remove them.
The other options are not correct commands for listing and removing corresponding packages. The dnf list command will display a list of available packages in enabled repositories, but not the packages installed by dnf transactions. The dnf remove command will remove specified packages from the system, but not all packages from a specific transaction. The dnf info command will display detailed information about specified packages, but not about dnf transactions. The dnf upgrade command will upgrade all installed packages to their latest versions, but not undo any package changes. References: Handling package management history; dnf(8) - Linux manual page

NEW QUESTION 226
A systems administrator is notified that the mysqld process stopped unexpectedly. The systems administrator issues the following command: sudo grep –i -r ‘out of memory’ /var/log
The output of the command shows the following:
kernel: Out of memory: Kill process 9112 (mysqld) score 511 or sacrifice child.
Which of the following commands should the systems administrator execute NEXT to troubleshoot this issue? (Select two).

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

Answer: AF

Explanation:
The free -h command can be used to check the amount of free and used memory in the system in a human-readable format. This can help to troubleshoot the issue of mysqld being killed due to out of memory. The vmstat -a 1 4 command can be used to monitor the system’s virtual memory statistics, such as swap usage, paging activity, and memory faults, every one second for four times. This can help to identify any memory pressure or performance issues that may cause out of memory errors. The nc -v 127.0.0.1 3306 command would attempt to connect to the MySQL server on port 3306 and display any diagnostic messages, but this would not help to troubleshoot the memory issue. The renice -15 $(pidof mysql) command would change the priority of the mysql process to -15, but this would not prevent it from being killed due to out of memory. The lsblk command would display information about block devices, not memory usage. The killall -15 command would send a SIGTERM signal to all processes with a matching name, but this would not help to troubleshoot the memory issue. References: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 15: Managing Memory and Process Execution, pages 468-469.

NEW QUESTION 228
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your XK0-005 Exam with Our Prep Materials Via below:

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