

Exam Questions XK0-005

CompTIA Linux+ Certification Exam

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



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 user is trying to execute commands with sudo but is receiving the following error:

```
$ sudo visudo
```

```
>>> /etc/sudoers: syntax error near line 28 <<< sudo: parse error in /etc/sudoers near line 28 sudo: no valid sudoers sources found, quitting The following output is provided:
```

```
# grep root /etc/shadow root :* LOCK *: 14600 :::::
```

Which of the following actions will resolve this issue?

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

Answer: B

NEW QUESTION 3

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

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

Answer: D

Explanation:

The command that will provide the correct information about whether files from the wget package have been altered since they were installed is rpm -V wget. This command will use the rpm utility to verify an installed RPM package by comparing information about the installed files with information from the RPM database. The verification process can check various attributes of each file, such as size, mode, owner, group, checksum, capabilities, and so on. If any discrepancies are found, rpm will report them using a single letter code for each attribute.

The other options are not correct commands for verifying an installed RPM package. The rpm -i wget command is invalid because -i is used to install a package from a file, not to verify an installed package. The rpm -qf wget command will query which package owns wget as a file name or path name, but it will not verify its attributes. The rpm -F wget command will freshen (upgrade) an already installed package with wget as a file name or path name, but it will not verify its attributes.

References: rpm(8) - Linux manual

page; Using RPM to Verify Installed Packages

NEW QUESTION 4

A 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 5

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

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

Answer: D

Explanation:

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

NEW QUESTION 6

A Linux administrator has physically added a new RAID adapter to a system. Which of the following commands should the Linux administrator run to confirm that the device has been recognized? (Select TWO).

- A. rmmod
- B. ls -l /etc
- C. lshw --class disk
- D. pvdisplay
- E. rmdir /dev
- F. dmesg

Answer: CF

Explanation:

The following commands can help you confirm that the new RAID adapter has been recognized by the Linux system:

? `dmesg`: This command displays the kernel messages, which can show the information about the newly detected hardware device. You can use `dmesg | grep -i raid` to filter the output for RAID-related messages.

? `lshw -class disk`: This command lists the disk devices on the system, including the RAID controller and its model name. You can use `lshw -class disk | grep -i raid` to filter the output for RAID-related information.

The other commands are not relevant for this purpose. For example:

? `rmmod`: This command removes a module from the Linux kernel, which is not useful for detecting a new device.

? `ls -l /etc`: This command lists the files and directories in the /etc directory, which is not related to hardware devices.

? `pvdisplay`: This command displays the attributes of physical volumes, which are part of the logical volume management (LVM) system, not the RAID system.

? rmdir /dev: This command removes an empty directory, which is not helpful for detecting a new device. Moreover, /dev is a special directory that contains device files, and should not be removed.

NEW QUESTION 7

A Linux administrator is configuring a new internal web server fleet. The web servers are up and running but can only be reached by users directly via IP address. The administrator is attempting to fix this inconvenience by requesting appropriate records from the DNS team. The details are:

Hostname: devel.comptia.org

IP address: 5.5.5.1, 5.5.5.2, 5.5.5.3, 5.5.5.4

Name server: 5.5.5.254

Additional names: dev.comptia.org, development.comptia.org

Which of the following types of DNS records should the Linux administrator request from the DNS team? (Select three).

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

Answer: BDE

Explanation:

The Linux administrator should request the following types of DNS records from the DNS team:

? A: This record type is used to map a hostname to an IPv4 address. The administrator needs four A records for devel.comptia.org, one for each IP address (5.5.5.1, 5.5.5.2, 5.5.5.3, 5.5.5.4). This will allow users to access the web servers by using the hostname devel.comptia.org instead of the IP addresses¹.

? CNAME: This record type is used to create an alias for another hostname. The administrator needs two CNAME records, one for dev.comptia.org and one for development.comptia.org, both pointing to devel.comptia.org. This will allow users to access the web servers by using any of these three hostnames interchangeably¹.

? NS: This record type is used to delegate a domain or a subdomain to another name server. The administrator needs one NS record for comptia.org, pointing to 5.5.5.254, which is the name server that hosts the records for the subdomain devel.comptia.org². This will allow users to resolve the hostnames under comptia.org by querying the name server 5.5.5.254.

The other record types are not relevant for the administrator's task:

? MX: This record type is used to specify the mail exchange server for a domain or a subdomain¹. The administrator does not need this record type because the web servers are not intended to handle email traffic.

? PTR: This record type is used to map an IP address to a hostname, which is the reverse of an A record¹. The administrator does not need this record type because the web servers are not expected to be accessed by their IP addresses.

? RRSIG: This record type is used to provide digital signatures for DNSSEC, which is a security extension for DNS that verifies the authenticity and integrity of DNS responses³. The administrator does not need this record type because it is not mentioned in the task requirements.

? SOA: This record type is used to provide information about the authoritative name server and other parameters for a domain or a subdomain¹. The administrator does not need this record type because it is usually created automatically by the name server software when a new zone file is created⁴.

? TXT: This record type is used to store arbitrary text data that can be used for various purposes, such as SPF, DKIM, DMARC, etc¹. The administrator does not need this record type because it is not related to the web server functionality.

? SRV: This record type is used to specify the location and port number of a service that runs on a domain or a subdomain¹. The administrator does not need this record type because the web servers use the standard HTTP port 80, which does not require an SRV record.

References: 1: DNS Record Types – CompTIA Network+ N10-007 – 1.8 2: NS Record - DNSimple Help 3: DNSSEC - Wikipedia 4: SOA Record - DNSimple Help

NEW QUESTION 8

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

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

Answer: D

Explanation:

This will prevent the file app.conf from being tracked by Git and uploaded to the repository. The .gitignore file is a special file that contains patterns of files and directories that Git should ignore. Any file that matches a pattern in the .gitignore file will not be staged, committed, or pushed to the remote repository. The .gitignore file should be placed in the root directory of the repository and committed along with the other files.

The other options are incorrect because:

* A. Run git exclude app.conf

This is not a valid Git command. There is no such thing as git exclude. The closest thing is git update-index --assume-unchanged, which tells Git to temporarily ignore changes to a file, but it does not prevent the file from being uploaded to the repository.

* B. Run git stash app.conf

This will temporarily save the changes to the file app.conf in a stash, which is a hidden storage area for uncommitted changes. However, this does not prevent the file from being tracked by Git or uploaded to the repository. The file will still be part of the working tree and the index, and it will be restored when the stash is popped or applied.

* C. Add app.conf to .exclude

This will have no effect, because Git does not recognize a file named .exclude. The only files that Git uses to ignore files are .gitignore, \$GIT_DIR/info/exclude, and core.excludesFile.

References:

? Git - gitignore Documentation

? .gitignore file - ignoring files in Git | Atlassian Git Tutorial

? Ignoring files - GitHub Docs

? [CompTIA Linux+ Certification Exam Objectives]

NEW QUESTION 9

A systems administrator wants to permit access temporarily to an application running on port 1234/TCP on a Linux server. Which of the following commands will permit this traffic?

- A. firewall-cmd --new-service=1234/tcp
- B. firewall-cmd --service=1234 --protocol=tcp
- C. firewall-cmd --add-port=1234/tcp
- D. firewall-cmd --add-whitelist-uid=1234

Answer: C

Explanation:

The firewall-cmd command is used to manage firewalld, which is a firewall service for Linux systems that provides dynamic and persistent configuration of firewall rules. Firewalld uses zones and services to define different levels of trust and access for network connections.

To permit access temporarily to an application running on port 1234/TCP on a Linux server, the systems administrator can use the firewall-cmd --add-port=1234/tcp command. This command will add a rule to the default zone (usually public) that allows incoming traffic on port 1234/TCP. The rule will only be effective until the next reload or restart of firewalld. To make the rule permanent, the administrator can add the --permanent option to the command. The statement C is correct.

The statements A, B, and D are incorrect because they do not permit access to port 1234/TCP. The firewall-cmd --new-service=1234/tcp command does not exist. The firewall-cmd --service=1234 --protocol=tcp command does not work because 1234 is not a predefined service name in firewalld. The firewall-cmd --add-whitelist-uid=1234 command does not exist. References: [How to Use Firewalld to Manage Firewall in Linux]

NEW QUESTION 10

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

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

# ethtool -S eth0
rx_fw_discards: 4487
```

Which of the following commands will improve the latency issue?

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

Answer: D

Explanation:

The best command to use to improve the latency issue is D. # echo 'net.core.rmem_max = 12500000' >> /etc/sysctl.conf # echo 'net.core.wmem_max = 12500000' >> /etc/sysctl.conf # sysctl -p. This command will increase the size of the receive and send buffers for the network interface, which can improve the network performance and reduce packet loss. The sysctl command will apply the changes to the kernel parameters without rebooting the system.

The other commands are either incorrect or not suitable for this task. For example:

? A. # echo 'net.core.net_backlog = 5000000' >> /etc/sysctl.conf # sysctl -p # systemctl daemon-reload will try to increase the backlog queue for incoming connections, but this is not relevant for the latency issue. The systemctl daemon-reload command is also unnecessary, as it only reloads the systemd configuration files, not the kernel parameters.

? B. # ifdown eth0 # ip link set dev eth0 mtu 800 # ifup eth0 will try to change the maximum transmission unit (MTU) of the network interface to 800 bytes, but this is too low and may cause fragmentation and performance degradation. The default MTU for Ethernet is 1500 bytes, and it should not be changed unless there is a specific reason.

? C. # systemctl stop network # ethtool -g eth0 512 # systemctl start network will try to change the ring buffer size of the network interface to 512, but this is too small and may cause packet drops and latency spikes. The default ring buffer size for Ethernet is usually 4096 or higher, and it should be increased if there is a high network traffic.

NEW QUESTION 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.sh chmod 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 12

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

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

Answer: D

Explanation:

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

NEW QUESTION 17

A user is unable to remotely log on to a server using the server name `server1` and port 22.

The Linux engineer troubleshoots the issue and gathers the following information: Which of the following is most likely causing the issue?

- A. `server 1` is not in the DNS.
- B. `sshd` is running on a non-standard port.
- C. `sshd` is not an active service.
- D. `server1` is using an incorrect IP address.

Answer: B

Explanation:

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

NEW QUESTION 19

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

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

Answer: C

Explanation:

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

NEW QUESTION 24

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

- A. `file filename`
- B. `touch filename`

- C. grep filename
- D. lsof filename

Answer: A

Explanation:

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

NEW QUESTION 25

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 26

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 28

A junior administrator is trying to set up a passwordless SSH connection to one of the servers. The administrator follows the instructions and puts the key in the authorized_key file at the server, but the administrator is still asked to provide a password during the connection.

Given the following output:

```
junior@server:~$ ls -lh .ssh/auth*
-rw----- 1 junior junior 566 sep 13 20:56 .ssh/authorized_key
```

Which of the following commands would resolve the issue and allow an SSH connection to be established without a password?

- A. restorecon -rv .ssh/authorized_key

- B. mv .ssh/authorized_key .ssh/authorized_keys
- C. systemctl restart sshd.service
- D. chmod 600 mv .ssh/authorized_key

Answer: B

Explanation:

The command mv .ssh/authorized_key .ssh/authorized_keys will resolve the issue and allow an SSH connection to be established without a password. The issue is caused by the incorrect file name of the authorized key file on the server. The file should be named authorized_keys, not authorized_key. The mv command will rename the file and fix the issue. The other options are incorrect because they either do not affect the file name (restorecon or chmod) or do not restart the SSH service (systemctl). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 410.

NEW QUESTION 33

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

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 37

A developer needs to launch an Nginx image container, name it Web001, and expose 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 Web001 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 40

Which of the following is the best tool for dynamic tuning of kernel parameters?

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

Answer: A

Explanation:

The tuned application is the best tool for dynamic tuning of kernel parameters, as it monitors the system and optimizes the performance under different workloads. It provides a number of predefined profiles for typical use cases, such as power saving, low latency, high throughput, virtual machine performance, and so on. It also allows users to create, modify, and delete profiles, and to switch between them on the fly. The tuned application uses the sysctl command and the configuration files in the /etc/sysctl.d/ directory to adjust the kernel parameters at runtime.

References

? Chapter 2. Getting started with TuneD - Red Hat Customer Portal, paragraph 1

? Kernel tuning with sysctl - Linux.com, paragraph 1

NEW QUESTION 43

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

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

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

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

Answer: CDH

Explanation:

To enable LUKS on a USB storage device with an ext4 filesystem format, the administrator needs to perform the following steps:

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

NEW QUESTION 44

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 49

Which of the following actions are considered good security practices when hardening a Linux server? (Select two).

- A. Renaming the root account to something else
- B. Removing unnecessary packages
- C. Changing the default shell to /bin/csh
- D. Disabling public key authentication
- E. Disabling the SSH root login possibility
- F. Changing the permissions on the root filesystem to 600

Answer: BE

Explanation:

Some good security practices when hardening a Linux server are:

- ? Removing unnecessary packages (B) to reduce the attack surface and eliminate potential vulnerabilities
 - ? Disabling the SSH root login possibility (E) to prevent unauthorized access and brute-force attacks on the root account
- References:
? [CompTIA Linux+ Study Guide], Chapter 9: Securing Linux, Section: Hardening Linux
? [How to Harden Your Linux Server]

NEW QUESTION 54

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

- A. clone
- B. gitignore
- C. get
- D. .ssh

Answer: B

Explanation:

To prevent certain files from being tracked by Git, the administrator can use a .gitignore file (B) in the repository. The .gitignore file can specify patterns of files or directories that Git should ignore. This way, the Terraform state files will not appear in the main branch or any other branch. The other commands are not related to this requirement. References:

? [CompTIA Linux+ Study Guide], Chapter 10: Working with Git, Section: Ignoring Files with .gitignore
 ? [How to Use .gitignore File]

NEW QUESTION 59

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

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

Answer: A

Explanation:

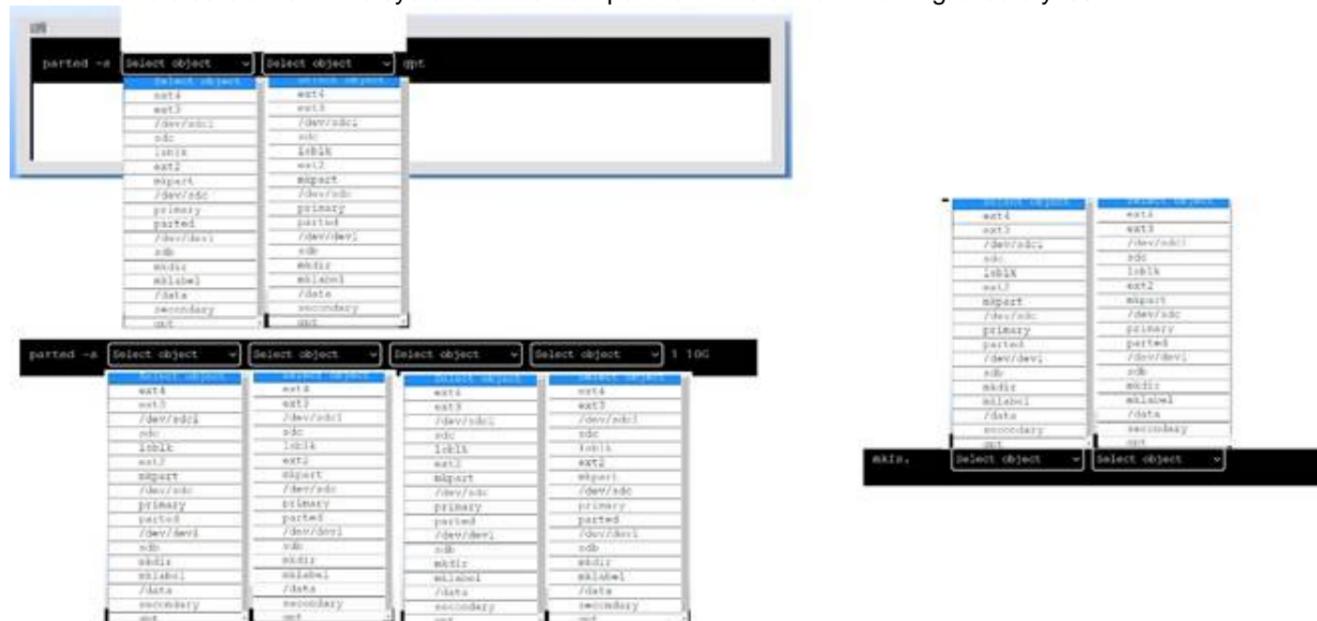
The command `systemctl status systemd-resolved.service` will show the information about the service `systemd-resolved.service`. The `systemctl` command is a tool for managing system services and units. The `status` option displays the current status of a unit, such as active, inactive, or failed. The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

NEW QUESTION 60

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 62

A user created the following script file:

```
#!/bin/bash
# FILENAME: /home/user/ script . sh echo "hello world"
exit 1
```

However, when the user tried to run the script file using the command `script . sh`, an error returned indicating permission was denied. Which of the following should the user execute in order for the script to run properly?

- A. `chmod u+x /home/user/script . sh`
- B. `chmod 600 /home/user/script . sh`
- C. `chmod /home/user/script . sh`

- D. `chmod 0+r /home/user/scrip`
- E. `sh`

Answer: A

Explanation:

To run a script file, the user needs to have execute permission on the file. The command `chmod u+x /home/user/script.sh` (A) will grant execute permission to the owner of the file, which is the user who created it. The other commands will not give execute permission to the user, and therefore will not allow the script to run properly. References:

- ? [CompTIA Linux+ Study Guide], Chapter 3: Working with Files, Section: Changing File Permissions
- ? [How to Make a Bash Script Executable]

NEW QUESTION 65

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 70

Due to low disk space, a Linux administrator finding and removing all log files that were modified more than 180 days ago. Which of the following commands will accomplish this task?

- A. `find /var/log -type d -mtime +180 -print -exec rm {} \;`
- B. `find /var/log -type f -modified +180 -rm`
- C. `find /var/log -type f -mtime +180 -exec rm {} \`
- D. `find /var/log -type c -atime +180 -remove`

Answer: C

Explanation:

The command that will accomplish the task of finding and removing all log files that were modified more than 180 days ago is `find /var/log -type f -mtime +180 -exec rm {} ;`. This command will use `find` to search for files (`-type f`) under `/var/log` directory that have a modification time (`-mtime`) older than 180 days (`+180`). For each matching file, it will execute (`-exec`) the `rm` command to delete it, passing the file name as an argument (`{}`). The command will end with a semicolon (`;`), which is escaped with a backslash to prevent shell interpretation.

The other options are not correct commands for accomplishing the task. The `find /var/log -type d -mtime +180 -print -exec rm {} ;` command will search for directories (`-type d`) instead of files, and print their names (`-print`) before deleting them. This is not what the task requires. The `find /var/log -type f -modified +180 -rm` command is invalid because there is no such option as `-modified` or `-rm` for `find`. The correct options are `-mtime` and `-delete`, respectively. The `find /var/log -type c -atime +180 -remove` command is also invalid because there is no such option as `-remove` for `find`. Moreover, it will search for character special files (`-type c`) instead of regular files, and use access time (`-atime`) instead of modification time. References: `find(1)` - Linux manual page; Find and delete files older than n days in Linux

NEW QUESTION 74

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

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

Answer: A

Explanation:

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

NEW QUESTION 75

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

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

The Linux server has the following system properties CPU: 4 vCPU
Memory: 50GB

Which of the following accurately describes this situation?

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

Answer: A

Explanation:

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

NEW QUESTION 76

A Linux administrator has set up a new DNS forwarder and is configuring all internal servers to use the new forwarder to look up external DNS requests. The administrator needs to modify the firewall on the server for the DNS forwarder to allow the internal servers to communicate to it and make the changes persistent between server reboots. Which of the following commands should be run on the DNS forwarder server to accomplish this task?

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

Answer: D

Explanation:

The command that should be run on the DNS forwarder server to accomplish the task is firewall-cmd --zone=public --add-port=53/udp --permanent.

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

NEW QUESTION 79

A systems administrator was tasked with assigning the temporary IP address/netmask 192.168.168.1/255.255.255.255 to the interface eth0 of a Linux server. When adding the address, the following error appears:

```
# ip address add 192.168.168.1/33 dev eth0
```

```
Error: any valid prefix is expected rather than "192.168.168.1/33".
```

Based on the command and its output above, which of the following is the cause of the issue?

- A. The CIDR value /33 should be /32 instead.
- B. There is no route to 192.168.168.1/33.
- C. The interface eth0 does not exist.
- D. The IP address 192.168.168.1 is already in use.

Answer: A

Explanation:

The cause of the issue is that the CIDR value /33 is invalid for an IPv4 address. The CIDR value represents the number of bits in the network prefix of an IP address, and it can range from 0 to 32 for IPv4 addresses. A CIDR value of /33 would imply a network prefix of more than 32 bits, which is impossible for an IPv4 address. To assign a temporary IP address/netmask of 192.168.168.1/255.255.255.255 to eth0, the CIDR value should be /32 instead, which means a network prefix of 32 bits and a host prefix of 0 bits. There is no route to 192.168.168.1/33 is not the cause of the issue, as the ip address add command does not check the routing table. The interface eth0 does not exist is not the cause of the issue, as the ip address add command would display a different error message if the interface does not exist. The IP address 192.168.168.1 is already in use is not the cause of the issue, as the ip address add command would display a different error message if the IP address is already in use. References: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 13: Networking Fundamentals, page 435.

NEW QUESTION 83

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 85

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

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

Answer: B

Explanation:

The command `docker cp container_id:/deployment.yaml deployment.yaml` can accomplish the task of copying the file `deployment.yaml` from a container to the host.

The `docker` command is a tool for managing Docker containers and images. The `cp` option copies files or directories between a container and the local filesystem. The `container_id` is the identifier of the container, which can be obtained by using the `docker ps` command.

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

NEW QUESTION 86

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

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

Answer: B

Explanation:

The best tool to use when uploading the public key to the remote servers is

* B. `ssh-copy-id`. This tool will copy the public key from the local computer to the remote server and append it to the `authorized_keys` file, which is used for public key authentication. This tool will also create the necessary directories and files on the remote server if they do not exist. The other tools are either not suitable or not relevant for this task. For example:

? A. `scp` is a tool for securely copying files between hosts, but it does not automatically add the public key to the `authorized_keys` file.

? C. `ssh-agent` is a tool for managing private keys and passphrases, but it does not upload the public key to the remote server.

? D. `ssh-keyscan` is a tool for collecting public keys from remote hosts, but it does not upload the public key to the remote server.

NEW QUESTION 91

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

Output 1:

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

Output 2:

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 93

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

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

Answer: B

Explanation:

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

NEW QUESTION 95

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

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

Answer: D

Explanation:

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

NEW QUESTION 99

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 | xargs grep -ln denied`
- B. `find . -type f -print | xargs grep -nv denied`
- C. `find . -type f -print | xargs grep -wL denied`
- D. `find . -type f -print | xargs 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 102

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 107

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 110

A developer wants to ensure that all files and folders created inside a shared folder named `/GroupOODEV` inherit the group name of the parent folder. Which of the following commands will help achieve this goal?

- A. `chmod g+X / GroupOODEV/`
- B. `chmod g+W / GroupOODEV/`
- C. `chmod g+r / GroupOODEV/`
- D. `chmod g+s / GroupOODEV/`

Answer: D

Explanation:

The `chmod` command is used to change the permissions of files and directories on Linux systems. The `g+s` option sets the setgid bit on a directory, which means that all files and folders created inside that directory will inherit the group name of the parent directory. This command can help the developer ensure that all files and folders created inside the `/GroupOODEV` directory have the same group name as `/GroupOODEV`. References: [How to Use `chmod` Command in Linux with Examples]

NEW QUESTION 113

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. `systemctl reboot`
- C. `systemctl get-default`
- D. `systemctl emergency`

Answer: B

Explanation:

The systemctl 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 systemctl get-default command would display the default target, not change it. The systemctl 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 115

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 116

Which of the following technologies provides load balancing, encryption, and observability in containerized environments?

- A. Virtual private network
- B. Sidecar pod
- C. Overlay network
- D. Service mesh

Answer: D

Explanation:

"A service mesh controls the delivery of service requests in an application. Common features provided by a service mesh include service discovery, load balancing, encryption and failure recovery."

The technology that provides load balancing, encryption, and observability in containerized environments is service mesh. A service mesh is a dedicated infrastructure layer that manages the communication and security between microservices in a distributed system. A service mesh consists of two components: a data plane and a control plane. The data plane is composed of proxies that are deployed alongside the microservices as sidecar pods. The proxies handle the network traffic between the microservices and provide features such as load balancing, encryption, authentication, authorization, routing, and observability. The control plane is responsible for configuring and managing the data plane and providing a unified interface for the administrators and developers. A service mesh can help improve the performance, reliability, and security of containerized applications and simplify the development and deployment process. A service mesh is the technology that provides load balancing, encryption, and observability in containerized environments. This is the correct answer to the question. The other options are incorrect because they either do not provide all the features of a service mesh (virtual private network or overlay network) or are not a technology but a component of a service mesh (sidecar pod). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 574. <https://www.techtarget.com/searchitoperations/definition/service-mesh>

NEW QUESTION 121

A systems administrator is implementing a new service task with systems at startup and needs to execute a script entitled test.sh with the following content:

```
TIMESTAMP=$(date '+%Y-%m-%d %H:%M:%S')
echo "helpme.service: timestamp $(Timestamp)" | systemd-cat -p info
sleep 60
done
```

The administrator tries to run the script after making it executable with chmod +x; however, the script will not run. Which of the following should the administrator do to address this issue? (Choose two.)

- A. Add #!/bin/bash to the bottom of the script.
- B. Create a unit file for the new service in /etc/systemd/system/ with the name helpme.service in the location.
- C. Add #!/bin/bash to the top of the script.
- D. Restart the computer to enable the new service.
- E. Create a unit file for the new service in /etc/init.d with the name helpme.service in the location.
- F. Shut down the computer to enable the new service.

Answer: BC

Explanation:

The administrator should do the following two things to address the issue:

? Add #!/bin/bash to the top of the script. This is called a shebang line and it tells the system which interpreter to use to execute the script. Without this line, the script will not run properly. The shebang line should be the first line of the script and should start with #! followed by the path to the interpreter. In this case, the interpreter is bash and the path is /bin/bash. The other option (A) is incorrect because the shebang line should be at the top, not the bottom of the script.

? Create a unit file for the new service in /etc/systemd/system/ with the name helpme.service in the location. This is necessary to register the script as a systemd service and enable it to run at startup. A unit file is a configuration file that defines the properties and behavior of a service, such as the description, dependencies,

start and stop commands, and environment variables. The unit file should have the extension `.service` and should be placed in the `/etc/systemd/system/` directory. The other option (E) is incorrect because `/etc/init.d` is the directory for init scripts, not systemd services.

References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, pages 429-430.

NEW QUESTION 123

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

- A. `[www1] rsync -a -e ssh /var/www/html/ user1@www2 : /var/www/html`
- B. `[www1] scp -r /var/www/html user1@www2 : /var/www/html`
- C. `[www2] cd /var/www/html; wget -m http://www1/`
- D. `[www1] cd /var/www/html && tar cvf -`

Answer: A

Explanation:

To mirror the website files on the primary web server, `www1`, to the backup web server, `www2`, the administrator can use the command `rsync -a -e ssh /var/www/html/ user1@www2:/var/www/html` (A). This will synchronize all files and directories under `/var/www/html/` on `www1` to `/var/www/html` on `www2` using `ssh` as the remote shell. The `-a` option will preserve all attributes and permissions of the files. The other commands will not mirror the website files, but either copy them once, download them from a web server, or archive them. References:

? [CompTIA Linux+ Study Guide], Chapter 12: Troubleshooting Linux Systems, Section: Synchronizing Files with `rsync`

? [How to Use `rsync` Command in Linux]

NEW QUESTION 127

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 129

A DevOps engineer is working on a local copy of a Git repository. The engineer would like to switch from the main branch to the staging branch but notices the staging branch does not exist. Which of the following Git commands should the engineer use to perform this task?

- A. `git branch -m staging`
- B. `git commit -m staging`
- C. `git status -b staging`
- D. `git checkout -b staging`

Answer: D

Explanation:

The correct answer is D. `git checkout -b staging`

This command will create a new branch named `staging` and switch to it. The `git checkout` command is used to switch between branches or restore files from a specific branch. The `-b` option is used to create a new branch if it does not exist. For example, `git checkout -b staging` will create and switch to the `staging` branch.

The other options are incorrect because:

* A. `git branch -m staging`

This command will rename the current branch to `staging`, not switch to it. The `git branch` command is used to list, create, or delete branches. The `-m` option is used to rename a branch. For example, `git branch -m staging` will rename the current branch to `staging`.

* B. `git commit -m staging`

This command will commit the changes in the working tree to the current branch with a message of `staging`, not switch to it. The `git commit` command is used to record changes to the repository. The `-m` option is used to specify a commit message. For example, `git commit -m staging` will commit the changes with a message of `staging`.

* C. `git status -b staging`

This command will show the status of the working tree and the current branch, not switch to it. The `git status` command is used to show the state of the working tree and the staged changes. The `-b` option is used to show the name of the current branch. However, this option does not take an argument, so specifying `staging` after it will cause an error. References:

? Git - `git-checkout` Documentation

? Git Tutorial: Create a New Branch With Git Checkout

? Git Branching - Basic Branching and Merging

NEW QUESTION 133

A database administrator requested the installation of a custom database on one of the servers. Which of the following should the Linux administrator configure so the requested packages can be installed?

- A. /etc/yum.conf
- B. /etc/ssh/sshd.conf
- C. /etc/yum.repos.d/db.repo
- D. /etc/resolv.conf

Answer: C

Explanation:

The Linux administrator should configure /etc/yum.repos.d/db.repo so that the requested packages can be installed. This file defines a custom repository for yum, which is a package manager for RPM-based systems. The file should contain information such as the name, baseurl, gpgcheck, and enabled options for the repository. By creating this file and enabling the repository, the administrator can use yum to install packages from the custom repository. The /etc/yum.conf file is the main configuration file for yum, but it does not define repositories. The /etc/ssh/sshd.conf file is the configuration file for sshd, which is a daemon that provides secure shell access to remote systems. The /etc/resolv.conf file is the configuration file for DNS resolution, which maps domain names to IP addresses. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Packages and Software, page 559.

NEW QUESTION 138

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

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

Answer: D

Explanation:

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

NEW QUESTION 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 143

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

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

Answer: CF

Explanation:

The administrator can use the following two options to boot the system into the single user mode:

? Interrupt the boot process in the GRUB menu and add systemd.unit=rescue.target in the kernel line. This option will boot the system into the rescue mode, which is a minimal environment that allows the administrator to perform basic tasks such as repairing the system. The GRUB menu is a screen that appears when the system is powered on and allows the administrator to choose which kernel or operating system to boot. The kernel line is a line that specifies the parameters for the kernel, such as the root device, the init system, and the boot options. The administrator can interrupt the boot process by pressing the e key in the GRUB menu and edit the kernel line by adding systemd.unit=rescue.target at the end. This option will tell the system to use the rescue target, which is a unit that defines the state of the system in the rescue mode. The administrator can then press Ctrl+X to boot the system with the modified kernel line. This option will boot the system into the single user mode and allow the administrator to troubleshoot the issues.

? Interrupt the boot process in the GRUB menu and add systemd.unit=single.target in the kernel line. This option will boot the system into the single user mode, which is a mode that allows the administrator to log in

as the root user and perform maintenance tasks. The GRUB menu and the kernel line are the same as the previous option. The administrator can interrupt the boot process by pressing the e key in the GRUB menu and edit the kernel line by adding systemd.unit=single.target at the end. This option will tell the system to use the single target, which is a unit that defines the state of the system in the single user mode. The administrator can then press Ctrl+X to boot the system with the modified kernel line. This option will boot the system into the single user mode and allow the administrator to troubleshoot the issues.

The other options are incorrect because they either do not boot the system into the single user mode (execute the following command from the GRUB rescue shell: mount -o remount, ro/sysroot or interrupt the boot process in the GRUB menu and add systemd.unit=single in the kernel line) or do not use the correct syntax (interrupt the boot process in the GRUB menu and add single=user in the kernel line or interrupt the boot process in the GRUB menu and add init=/bin/bash in the kernel

line). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing the Linux Boot Process, pages 267-268.

NEW QUESTION 148

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

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

Chain FORWARD (policy ACCEPT)
target      prot opt source      destination

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

The systems administrator makes additional checks:

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

firewalld is not running
```

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

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

Answer: D

Explanation:

The reason the firewall rules are not active is that the firewalld service is not enabled. This means that the service will not start automatically at boot time or after a system reload. To enable the firewalld service, the systems administrator needs to use the command sudo systemctl enable firewalld. This will create a symbolic link from the firewalld service file to the appropriate systemd target, such as multi-user.target. Enabling the service does not start it immediately, so the systems administrator also needs to use the command sudo systemctl start firewalld or sudo systemctl reload firewalld to activate the firewall rules.

The other options are not correct reasons for the firewall rules not being active. iptables is not conflicting with firewalld, because firewalld uses iptables as its backend by default. The wrong system target is not activated, because firewalld is independent of the system target and can be enabled for any target.

FIREWALL_ARGS has no value assigned, but this is not a problem, because FIREWALL_ARGS is an optional environment variable that can be used to pass additional arguments to the firewalld daemon, such as --debug or --nofork. If FIREWALL_ARGS is empty or not defined, firewalld will use its default arguments. References: firewalld.service(8) - Linux manual page; firewall-cmd(1) - Linux manual page; systemctl(1) - Linux manual page

NEW QUESTION 152

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. umask
- C. 002
- D. chattr -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 157

A Linux administrator needs to create a symlink for `/usr/local/bin/app-a`, which was installed in `/usr/local/share/app-a`. Which of the following commands should the administrator use?

- A. `ln -s /usr/local/bin/app-a /usr/local/share/app-a`
- B. `mv -f /usr/local/share/app-a /usr/local/bin/app-a`
- C. `cp -f /usr/local/share/app-a /usr/local/bin/app-a`
- D. `rsync -a /usr/local/share/app-a /usr/local/bin/app-a`

Answer: A

Explanation:

To create a symlink for `/usr/local/bin/app-a`, which was installed in `/usr/local/share/app-a`, the administrator can use the command `ln -s /usr/local/share/app-a /usr/local/bin/app-a` (A). This will create a symbolic link named `/usr/local/bin/app-a` that points to the original file `/usr/local/share/app-a`. The other commands will not create a symlink, but either move, copy, or synchronize the file. References:
? [CompTIA Linux+ Study Guide], Chapter 3: Working with Files, Section: Creating Links
? [How to Create Symbolic Links in Linux]

NEW QUESTION 158

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

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

Answer: B

Explanation:

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

NEW QUESTION 162

A new application container was built with an incorrect version number. Which of the following commands should be used to rename the image to match the correct version 2.1.2?

- A. `docker tag comptia/app:2.1.1 comptia/app:2.1.2`
- B. `docker push comptia/app:2.1.1 comptia/app:2.1.2`
- C. `docker rmi comptia/app:2.1.1 comptia/app:2.1.2`
- D. `docker update comptia/app:2.1.1 comptia/app:2.1.2`

Answer: A

Explanation:

The best command to use to rename the image to match the correct version 2.1.2 is A. `docker tag comptia/app:2.1.1 comptia/app:2.1.2`. This command will create a new tag for the existing image with the new version number, without changing the image content or ID. The other commands are either incorrect or not suitable for this task. For example:
? B. `docker push comptia/app:2.1.1 comptia/app:2.1.2` will try to push two images to a remote repository, but it does not rename the image locally.
? C. `docker rmi comptia/app:2.1.1 comptia/app:2.1.2` will try to remove two images from the local system, but it does not rename the image.
? D. `docker update comptia/app:2.1.1 comptia/app:2.1.2` will try to update the configuration of a running container, but it does not rename the image.

NEW QUESTION 164

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

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

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

[Install]
WantedBy=timers.target
```

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 168

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

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

Answer: A

Explanation:

The command chmod 755 filename.log should be used to set filename.log permissions to -rwxr--r--. The chmod command is a tool for changing file permissions on Linux file systems. The permissions can be specified in octal notation, where each digit represents the permissions for the owner, group, and others respectively. The permissions are encoded as follows:

- ? 0: no permission
- ? 1: execute permission
- ? 2: write permission
- ? 4: read permission
- ? 5: read and execute permissions (4 + 1)
- ? 6: read and write permissions (4 + 2)
- ? 7: read, write, and execute permissions (4 + 2 + 1)

The command chmod 755 filename.log will set the permissions to -rwxr--r--, which means that the owner has read, write, and execute permissions (7), the group has read and execute permissions (5), and others have read and execute permissions (5). This is the correct command to use to accomplish the task. The other options are incorrect because they either set the wrong permissions (chmod 640, chmod 740, or chmod 744) or do not exist (chmod -G). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, page 345.

NEW QUESTION 169

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 172

A Linux administrator booted up the server and was presented with a non-GUI terminal. The administrator ran the command `systemctl isolate graphical.target` and rebooted the system by running `systemctl 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 `systemctl 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 `systemctl 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 `systemctl 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: `systemctl(1)` - Linux manual page; How to Change Runlevels (targets) in SystemD

NEW QUESTION 174

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. `systemctl status firewalld`
- D. `firewall-cmd --list-all`
- 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 178

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

- A. `rpm -qf /bin/someapp`
- B. `rpm -Vv /bin/someapp`
- C. `rpm -P /bin/some app`
- D. `rpm -i /bin/someapp`

Answer: A

Explanation:

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

NEW QUESTION 179

Users are experiencing high latency when accessing a web application served by a Linux machine. A systems administrator checks the network interface counters

and sees the following:

```
# ip -s link list dev enp0s25
2: enp0s25: <BROADCAST,MULTICAST,LOWER_UP,UP> mtu 1500 qdisc fq_codel state DOWN mode DEFAULT group default qlen 1000 link/ether
ac:12:34:56:78:cd brd ff:ff:ff:ff:ff:ff
RX: bytes  packets  errors  dropped missed  mcast
2011664755 3579033 2394390 508    0    0
TX: bytes  packets  errors  dropped carrier collsns
309541780 1705408  0      0      12340  0
```

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

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

Answer: B

Explanation:

The high number of errors and dropped packets in the output of the network interface counters indicate a connection problem on the network interface.

References:

? CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Networking, Section: Troubleshooting Network Issues, Page 359.

? Linux+ (Plus) Certification, Exam Objectives: 4.3 Given a scenario, troubleshoot and resolve basic network configuration and connectivity issues.

NEW QUESTION 183

A Linux administrator is providing a new Nginx image from the registry to local cache. Which of the following commands would allow this to happen?

- A. docker pull nginx
- B. docker attach nginx
- C. docker commit nginx
- D. docker import nginx

Answer: A

Explanation:

The command that would allow this to happen is docker pull nginx. Docker is a software platform that allows the administrator to create, run, and manage containers on Linux systems. Containers are isolated and lightweight environments that can run applications and services without affecting the host system. Docker uses images to create containers, which are files that contain the code, libraries, dependencies, and configuration of the applications and services. Docker uses a registry to store and distribute images, which is a service that hosts and serves images. Docker Hub is the default public registry that provides a large number of official and community images. Nginx is a popular web server and reverse proxy that can run as a container. The command docker pull nginx will download the latest version of the Nginx image from the Docker Hub registry to the local cache, which is the storage location for the images on the host system. This will allow the administrator to provide a new Nginx image from the registry to the local cache. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not download an image from the registry (docker attach nginx or docker commit nginx) or do not exist (docker import nginx). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 571.

NEW QUESTION 186

As part of the requirements for installing a new application, the swappiness parameter needs to be changed to 0. This change needs to persist across re-boots and be applied immediately. A Linux systems administrator is performing this change. Which of the following steps should the administrator complete to accomplish this task?

- A. echo "v
- B. swappiness=0" >> /etc/sysctl.conf && sysctl -p
- C. echo "vr
- D. >> /proc/meminfo && sysctl -a
- E. sysctl -v >> /proc/meminfo && echo "v
- F. swappiness=0"
- G. sysctl -h "v
- H. swappiness=0" && echo /etc/vm.swappiness

Answer: A

Explanation:

To change the swappiness parameter to 0 and make it persistent across reboots and applied immediately, the administrator can perform the following steps:

? Append the line vm.swappiness=0 to the file /etc/sysctl.conf using echo

"vm.swappiness=0" >> /etc/sysctl.conf (A). This will set the swappiness parameter to 0 for future boots.

? Reload the sysctl configuration using sysctl -p (A). This will apply the changes to the current system without rebooting. The other commands will not achieve this task, but either write to a wrong file, use a wrong option, or have a syntax error. References:

? [CompTIA Linux+ Study Guide], Chapter 8: Optimizing Linux Performance, Section: Tuning Kernel Parameters with sysctl

? [How to Change Swappiness in Linux]

NEW QUESTION 190

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

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

Answer: B

Explanation:

Podman `exec -ti app /bin/sh` allows the Linux administrator to enter the running container and analyze the logs that are stored inside. This command uses the podman tool, which is a daemonless container engine that can run and manage containers on Linux systems. The `exec` option executes a command inside an existing container, in this case `app`, which is the name of the container that runs the failing application. The `-ti` option allocates a pseudo-TTY and keeps STDIN open, allowing for interactive shell access to the container. The `/bin/sh` argument specifies the shell command to run inside the container, which can be used to view and manipulate the log files.

The other options are not correct commands for entering a running container and analyzing the logs. `Docker run -ti app /bin/sh` creates a new container from the `app` image and runs the `/bin/sh` command inside it, but does not enter the existing container that runs the failing application. `Podman run -d app /bin/bash` also creates a new container from the `app` image and runs the `/bin/bash` command inside it, but does so in detached mode, meaning that it runs in the background without interactive shell access. `Docker exec -d app /bin/bash` executes the `/bin/bash` command inside the existing `app` container, but also does so in detached mode, without interactive shell access.

References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks; View container logs | Docker Docs; How to see the logs of a docker container - Stack Overflow

NEW QUESTION 191

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 196

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 200

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

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

Answer: A

Explanation:

The `su` command is used to switch to another user account on Linux systems. The `-` option makes the shell a login shell, which means that it will read the profile and environment variables of the target user. Without this option, the shell will retain the environment variables of the original user. This can cause confusion when issuing commands that depend on these variables, such as `ls`, which uses the `$HOME` variable to determine the home directory. Therefore, Ann should have issued `su - Joe` to list Joe's files instead of her own. References: [How to Use su Command in Linux with Examples]

NEW QUESTION 202

A Linux administrator has logged in to a server for the first time and needs to know which services are allowed through the firewall. Which of the following options will return the results for which the administrator is looking?

- A. firewall-cmd --get-services
- B. firewall-cmd --check-config
- C. firewall-cmd --list-services
- D. systemctl status firewalld

Answer: C

Explanation:

The firewall-cmd --list-services command will return the results for which the administrator is looking. This command will list all services that are allowed through the firewall in the default zone or a specified zone. A service is a predefined set of ports and protocols that can be enabled or disabled by firewalld. The firewall-cmd --get-services command will list all available services that are supported by firewalld, not only those that are allowed through the firewall. The firewall-cmd --check-config command will check if firewalld configuration files are valid, not list services. The systemctl status firewalld command will display information about the firewalld service unit, such as its state, PID, memory usage, and logs, not list services. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 543.

NEW QUESTION 204

An administrator attempts to rename a file on a server but receives the following error.

```
mv: cannot move 'files/readme.txt' to 'files/readme.txt.orig': operation not permitted.
```

The administrator then runs a few commands and obtains the following output:

```
$ ls -ld files/
drwxrwxrwt.1  users  users  20   Sep 10 15:15  files/
$ ls -a files/
drwxrwxrwt.1  users  users  20   Sep 10 15:15  -
drwxr-xr-x.1  users  users  32   Sep 10 15:15  ..
-rw-rw-r--.1  users  users   4   Sep 12 10:34  readme.txt
```

Which of the following commands should the administrator run NEXT to allow the file to be renamed by any user?

- A. chgrp reet files
- B. chacl -R 644 files
- C. chown users files
- D. chmod -t files

Answer: D

Explanation:

The command that the administrator should run NEXT to allow the file to be renamed by any user is chmod -t files. This command uses the chmod tool, which is used to change file permissions and access modes. The -t option removes (or sets) the sticky bit on a directory, which restricts deletion or renaming of files within that directory to only their owners or root. In this case, since files is a directory with sticky bit set (indicated by t in drwxrwxrwt), removing it will allow any user to rename or delete files within that directory. The other options are not correct commands for allowing any user to rename files within files directory. The chgrp reet files command will change the group ownership of files directory to reet, but it will not affect its permissions or access modes. The chacl -R 644 files command is invalid, as chacl is used to change file access control lists (ACLs), not permissions or access modes. The chown users files command will change the user ownership of files directory to users, but it will not affect its permissions or access modes. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing Users and Groups; chmod(1) - Linux manual page

NEW QUESTION 206

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

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

Answer: C

Explanation:

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

NEW QUESTION 208

A cloud engineer needs to launch a container named web-01 in background mode. Which of the following commands will accomplish this task?"

- A. docker builder -f --name web-01 httpd
- B. docker load --name web-01 httpd
- C. docker ps -a --name web-01 httpd
- D. docker run -d --name web-01 httpd

Answer: D

Explanation:

The docker run -d --name web-01 httpd command will launch a container named web-01 in background mode. This command will create and start a new container from the httpd image, assign it the name web-01, and run it in detached mode (-d), which means the container will run in the background without attaching to the current terminal. The docker builder -f --name web-01 httpd command is invalid, as builder is not a valid docker command, and -f and --name are not valid options for docker build. The docker load --name web-01 httpd command is invalid, as load does not accept a --name option, and httpd is not a valid file name for load. The docker ps -a --name web-01 httpd command is invalid, as ps does not accept a --name option, and httpd is not a valid filter for ps. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Virtualization and Cloud Technologies, page 499.

NEW QUESTION 210

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

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

Which of the following best explains the reported issue?

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

Answer: D

Explanation:

Based on the output from one of the affected systems, the best explanation for the reported issue is that the virtual machine is running out of CPU resources, leading to users experiencing longer response times (D). The output shows that the system use percentage is very high (57.85%), indicating that the virtual machine is using most of its CPU cycles for system processes. This leaves little CPU time for user processes, which results in slower performance. The other explanations are not supported by the output or are contradictory. References:

- ? [CompTIA Linux+ Study Guide], Chapter 8: Optimizing Linux Performance, Section: Monitoring CPU Usage
- ? [How to Interpret CPU Usage Statistics]

NEW QUESTION 214

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

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

Answer: A

Explanation:

The command that will accomplish the goal of avoiding the deletion or modification of files from non-owners is chmod +t /project/access2all. This command will set the sticky bit on the directory /project/access2all, which is a special permission that restricts file deletion or renaming to only the file owner, directory owner, or root user. This way, even if multiple users have write permission to the directory, they cannot delete or modify each other's files.

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

NEW QUESTION 219

A Linux administrator generated a list of users who have root-level command-line access to the Linux server to meet an audit requirement. The administrator analyzes the following /etc/passwd and /etc/sudoers files:

```
$ cat /etc/passwd
root:x:0:0:/home/root:/bin/bash lee:x:500:500:/home/lee:/bin/tcsh
mallory:x:501:501:/root:/bin/bash
eve:x:502:502:/home/eve:/bin/nologin carl:x:0:503:/home/carl:/bin/sh
bob:x:504:504:/home/bob:/bin/ksh
alice:x:505:505:/home/alice:/bin/rsh
$ cat /etc/sudoers
Cmnd_Alias SHELLS = /bin/tcsh, /bin/sh, /bin/bash Cmnd_Alias SYSADMIN = /usr/sbin/tcpdump
ALL = (ALL) ALL
```

ALL = NOPASSWD: SYSADMIN

Which of the following users, in addition to the root user, should be listed in the audit report as having root-level command-line access? (Select two).

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

Answer: AC

Explanation:

The users who have root-level command-line access are those who have either the same user ID (UID) as root, which is 0, or the ability to run commands as root using sudo. Based on the `/etc/passwd` and `/etc/sudoers` files, the users who meet these criteria are:

? Carl: Carl has the same UID as root, which is 0, as shown in the `/etc/passwd` file.

This means that Carl can log in as root and execute any command with root privileges¹

? Mallory: Mallory has the ability to run commands as root using sudo, as shown in the `/etc/sudoers` file. The line `ALL = (ALL) ALL` means that any user can run any command as any other user, including root, by using sudo. Mallory can also use the root shell `/bin/bash` as her login shell, as shown in the `/etc/passwd` file²

Therefore, the correct answer is A and C. Lee, Eve, Bob, and Alice do not have root-level command-line access because they have different UIDs from root and they cannot use sudo to run commands as root. Lee can only use sudo to run the commands listed in the `Cmnd_Alias SHELLS`, which are `/bin/tcsh`, `/bin/sh`, and `/bin/bash`. Eve cannot log in at all because her login shell is `/bin/nologin`. Bob and Alice can only use sudo to run the command `/usr/sbin/tcpdump` without a password, as specified by the `Cmnd_Alias SYSADMIN` and the line `ALL = NOPASSWD: SYSADMIN2`

NEW QUESTION 220

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

XK0-005 Practice Exam Features:

- * XK0-005 Questions and Answers Updated Frequently
- * XK0-005 Practice Questions Verified by Expert Senior Certified Staff
- * XK0-005 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * XK0-005 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year