

XK0-005 Dumps

CompTIA Linux+ Certification Exam

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



NEW QUESTION 1

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

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

Answer: A

Explanation:

This command will load the KVM module as well as any related dependencies, such as kvm-intel or kvm-amd, depending on the processor type. The modprobe command is a Linux utility that reads the /etc/modules.conf file and adds or removes modules from the kernel. It also resolves any dependencies between modules, so that they are loaded in the correct order.

The other options are incorrect because:

* B. insmod kvm

This command will only load the KVM module, but not any related dependencies. The insmod command is a low-level Linux utility that inserts a single module into the kernel. It does not resolve any dependencies between modules, so they have to be loaded manually.

* C. depmod kvm

This command will not load the KVM module at all, but only create a list of module dependencies for modprobe to use. The depmod command is a Linux utility that scans the installed modules and generates a file called modules.dep that contains dependency information for each module.

* D. hotplug kvm

This command is invalid and does not exist. The hotplug mechanism is a feature of the Linux kernel that allows devices to be added or removed while the system is running. It does not have anything to do with loading modules.

NEW QUESTION 2

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

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

Answer: D

Explanation:

The RPM option -e should be used to remove software from the server. The rpm command is a tool for managing software packages on RPM-based Linux distributions. The -e option stands for erase and removes the specified package from the system. This is the correct option to use to accomplish the task. The other options are incorrect because they either do not exist (-s or -d) or do not remove software (-q stands for query and displays information about the package).

References: CompTIA Linux+ (XK0-

005) Certification Study Guide, Chapter 16: Managing Software, page 489.

NEW QUESTION 3

A systems administrator received a notification that a system is performing slowly. When running the top command, the systems administrator can see the following values:

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 4

A user reported issues when trying to log in to a Linux server. The following outputs were received:

Given the outputs above, which of the following is the reason the user is unable to log in to the server?

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

Answer: D

Explanation:

The user1 password is expired. This can be inferred from the output of the `chage -l user1` command, which shows the password expiration information for user1. The output shows that the password expired on 2020-10-01, and the account expired on 2020-10-08. This means that user1 cannot log in to the server unless the password and account are reactivated by the system administrator.

The other options are not correct based on the outputs above. User1 does not need to set a long password, because the output of the `passwd -S user1` command shows that the password has a minimum length of 5 characters, which is met by user1's password. User1 is not in the incorrect group, because the output of the `groups user1` command shows that user1 belongs to the app group, which is presumably the correct group for accessing the server. The user1 shell assignment is not incorrect, because the output of the `grep user1 /etc/passwd` command shows that user1 has `/bin/bash` as the default shell, which is a valid and common shell for Linux users.

NEW QUESTION 5

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 6

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

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

Answer: A

Explanation:

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

NEW QUESTION 7

A non-privileged user is attempting to use commands that require elevated account permissions, but the commands are not successful. Which of the following most likely needs to be updated?

- A. `/etc/passwd`
- B. `/etc/shadow`
- C. `/etc/sudoers`
- D. `/etc/bashrc`

Answer: C

Explanation:

The `/etc/sudoers` file is used to configure the `sudo` command, which allows non-privileged users to execute commands that require elevated account permissions¹. The file contains a list of users and groups that are allowed to use `sudo`, and the commands they can run with it. The file also defines the security policy for `sudo`, such as whether a password is required, how long the `sudo` session lasts, and what environment variables are preserved or reset.

The `/etc/passwd` file is used to store information about the user accounts on the system, such as their username, user ID, home directory, and login shell. The `/etc/shadow` file is used to store the encrypted passwords for the user accounts, along with other information such as password expiration and aging. These files are not directly related to the `sudo` command, and updating them will not grant a user elevated account permissions.

The `/etc/bashrc` file is used to set up the environment for the `bash` shell, such as aliases, functions, variables, and options. This file is executed whenever a new `bash` shell is started, and it affects all users on the system. However, this file does not control the `sudo` command or its configuration, and updating it will not allow a user to use commands that require elevated account permissions.

NEW QUESTION 8

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

- A. `parted`
- B. `df`

- C. mount
- D. du
- E. fdisk
- F. dd
- G. ls

Answer: BD

Explanation:

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

NEW QUESTION 9

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

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

Answer: A

Explanation:

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

NEW QUESTION 10

A Linux administrator is troubleshooting a systemd mount unit file that is not working correctly. The file contains:

```
[root@system] # cat mydocs.mount [Unit]
Description=Mount point for My Documents drive [Mount]
What=/dev/drv/disk/by-uuid/94afc9b2-ac34-ccff-88ae-297ab3c7ff34 Where=/home/user1/My Documents
Options=defaults Type=xfs
[Install]
WantedBy=multi-user.target
```

The administrator verifies the drive UUID correct, and user1 confirms the drive should be mounted as My Documents in the home directory. Which of the following can the administrator do to fix the issues with mounting the drive? (Select two).

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

Answer: AE

Explanation:

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

NEW QUESTION 10

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

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

Answer: D

Explanation:

The command that will allow the technician to execute the services and continue deploying other microservices within the same terminal session is bg %1 job name. This command will send the job with ID 1 and name job name to the background, where it will run without occupying the terminal. The other options are incorrect because:

? gedit & disown will launch a graphical text editor in the background and detach it from the terminal, but it will not execute any service.
? kill 9 %1 will terminate the job with ID 1 using a SIGKILL signal, which cannot be ignored or handled by the process.
? fg %1 will bring the job with ID 1 to the foreground, where it will occupy the terminal until it finishes or is stopped. References: CompTIA Linux+ Study Guide, Fourth Edition, page 181-182.

NEW QUESTION 11

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

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

Answer: C

Explanation:

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

NEW QUESTION 13

To harden one of the servers, an administrator needs to remove the possibility of remote administrative login via the SSH service. Which of the following should the administrator do?

- A. Add the line `DenyUsers root` to the `/etc/hosts.deny` file.
- B. Set `PermitRootLogin` to `no` in the `/etc/ssh/sshd_config` file.
- C. Add the line `account required pam_nologin`
- D. so to the `/etc/pam.d/sshd` file.
- E. Set `PubKeyAuthentication` to `no` in the `/etc/ssh/ssh_config` file.

Answer: B

Explanation:

The administrator should set `PermitRootLogin` to `no` in the `/etc/ssh/sshd_config` file to remove the possibility of remote administrative login via the SSH service. The `PermitRootLogin` directive controls whether the root user can log in using SSH. Setting it to `no` will deny any remote login attempts by the root user. This will harden the server and prevent unauthorized access. The administrator should also restart the `sshd` service after making the change. The other options are incorrect because they either do not affect the SSH service (`/etc/hosts.deny` or `/etc/pam.d/sshd`) or do not prevent remote administrative login (`PubKeyAuthentication`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 413.

NEW QUESTION 18

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

- A. `systemctl isolate multi-user.target` `sh script.sh` `systemctl isolate graphical.target`
- B. `systemctl isolate graphical.target` `sh script.sh` `systemctl isolate multi-user.target`
- C. `sh script.sh` `systemctl isolate multi-user.target` `systemctl isolate graphical.target`
- D. `systemctl isolate multi-user.target` `systemctl isolate graphical.target` `sh script.sh`

Answer: A

Explanation:

The correct answer is A. `systemctl isolate multi-user.target` `sh script.sh` `systemctl isolate graphical.target`

This sequence will allow the administrator to switch from the graphical mode to the console mode, run the script, and then switch back to the graphical mode.

The `systemctl` command is used to control the `systemd` system and service manager, which manages the boot targets and services on Linux systems. The `isolate` subcommand starts the unit specified on the command line and its dependencies and stops all others. The `multi-user.target` is a boot target that provides a text-based console login, while the `graphical.target` is a boot target that provides a graphical user interface. By using `systemctl isolate`, the administrator can change the boot target on the fly without rebooting the system.

The `sh` command is used to run a shell script, which is a file that contains a series of commands that can be executed by the shell. The `script.sh` is the name of the script that contains the application change that the administrator needs to make. By running `sh script.sh`, the administrator can execute the script in the console mode.

The other options are incorrect because:

* B. `systemctl isolate graphical.target` `sh script.sh` `systemctl isolate multi-user.target`

This sequence will switch from the console mode to the graphical mode, run the script, and then switch back to the console mode. This is not what the administrator wants to do, as the script must be run only in console mode.

* C. `sh script.sh` `systemctl isolate multi-user.target` `systemctl isolate graphical.target`

This sequence will run the script in the current mode, which may or may not be console mode, and then switch to console mode and back to graphical mode. This is not what the administrator wants to do, as the script must be run only in console mode.

* D. `systemctl isolate multi-user.target` `systemctl isolate graphical.target` `sh script.sh`

This sequence will switch from graphical mode to console mode and then back to graphical mode, without running the script at all. This is not what the administrator wants to do, as the script must be run only in console mode.

References:

? `systemctl(1)` - Linux manual page

? How to switch between the CLI and GUI on a Linux server

? How to PROPERLY boot into single user mode in RHEL/CentOS 7/8

? Changing Systemd Boot Target in Linux

? Exit Desktop to Terminal in Ubuntu 19.10

NEW QUESTION 20

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

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

Answer: C

Explanation:

The ~/.ssh/config file can be used to set various options for SSH connections, including the port number, for specific hosts or groups of hosts. This file is located in the user's home directory and affects only the current user. The /etc/ssh/sshd_config file is used to configure the SSH server daemon, not the client. The /etc/ssh/moduli file contains parameters for Diffie-Hellman key exchange, not port settings.

The ~/.ssh/authorized_keys file contains public keys for authentication, not port settings. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Secure Shell (SSH), page 414.

NEW QUESTION 23

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

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

Answer: B

Explanation:

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

NEW QUESTION 24

An administrator runs ping comptia.org. The result of the command is:

ping: comptia.org: Name or service not known

Which of the following files should the administrator verify?

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

Answer: C

Explanation:

The best file to verify when the ping command returns the error "Name or service not known" is C. /etc/resolv.conf. This file contains the configuration for the DNS resolver, which is responsible for translating domain names into IP addresses. If this file is missing, corrupted, or has incorrect entries, the ping command will not be able to resolve the domain name and will fail with the error. To fix this issue, the administrator should check that the file exists, has proper permissions, and has valid nameserver entries. For example, a typical /etc/resolv.conf file may look like this:

```
nameserver 8.8.8.8 nameserver 8.8.4.4
```

These are the IP addresses of Google's public DNS servers, which can be used as a fallback option if the default DNS servers are not working.

NEW QUESTION 29

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 32

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

- A. visudo -c
- B. test -f /etc/sudoers
- C. sudo vi check

D. cat /etc/sudoers | tee test

Answer: A

Explanation:

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

NEW QUESTION 37

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

- A. The server's CPU is taking too long to process users' requests.
- B. The server's CPU shows a high idle-time value.
- C. The server's CPU is spending too much time waiting for data inputs.
- D. The server's CPU value for the time spent on system processes is low.

Answer: C

Explanation:

The server's CPU is spending too much time waiting for data inputs. This can be inferred from the output of the top command, which shows the percentage of CPU time spent in different states. The wa state stands for wait, and it indicates that the CPU is idle while waiting for an I/O operation to complete. In this case, the wa state is 52.4%, which means that more than half of the CPU time is wasted on waiting for data inputs. This can cause a high I/O latency and affect the performance of the server.

The other options are not correct based on the output received from the executed command. The server's CPU is not taking too long to process users' requests, because the us state, which stands for user, is only 0.2%, which means that the CPU is barely used by user processes. The server's CPU does not show a high idle-time value, because the id state, which stands for idle, is 0.0%, which means that the CPU is not idle at all. The server's CPU value for the time spent on system processes is not low, because the sy state, which stands for system, is 33.1%, which means that the CPU is heavily used by system processes. References: How to Use the Linux top Command (and Understand Its Output); [Understanding Linux CPU Load - when should you be worried?]

NEW QUESTION 42

A cloud engineer needs to change the secure remote login port from 22 to 49000. Which of the following files should the engineer modify to change the port number to the desired value?

- A. /etc/host.conf
- B. /etc/hostname
- C. /etc/services
- D. /etc/ssh/sshd_config

Answer: D

Explanation:

The file /etc/ssh/sshd_config contains the configuration settings for the SSH daemon, which handles the secure remote login. To change the port number, the engineer should edit this file and modify the line that says Port 22 to Port 49000. The other files are not related to the SSH service. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 411.

NEW QUESTION 44

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 45

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 49

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

Output 1:

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

Output 2:

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

Output 3:

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 51

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

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

Answer: ACE

Explanation:

The administrator should use the following three commands to resolve the issue of the root filesystem being full:

? df -h /. This command will show the disk usage of the root filesystem in a human-readable format. The df command is a tool for reporting file system disk space usage. The -h option displays the sizes in powers of 1024 (e.g., 1K, 234M, 2G). The / specifies the root filesystem. The command df -h / will show the total size, used space, available space, and percentage of the root filesystem. This command will help the administrator identify the problem and plan the solution.

? growpart /dev/mapper/rootvg-rootlv. This command will grow the partition that contains the root filesystem to the maximum size available.

The growpart command is a tool for resizing partitions on Linux systems. The /dev/mapper/rootvg-rootlv is the device name of the partition, which is a logical volume managed by the Logical Volume Manager (LVM). The command growpart /dev/mapper/rootvg-rootlv will extend the partition to fill the disk space and increase the size of the root filesystem. This command will help the administrator solve the problem and free up space.

? lvresize -L +10G -r /dev/mapper/rootvg-rootlv. This command will resize the logical volume that contains the root filesystem and add 10 GB of space.

The `lvresize` command is a tool for resizing logical volumes on Linux systems. The `-L` option specifies the new size of the logical volume, in this case `+10G`, which means 10 GB more than the current size. The `-r` option resizes the underlying file system as well. The `/dev/mapper/rootvg-rootlv` is the device name of the logical volume, which is the same as the partition name. The command `lvresize -L +10G -r /dev/mapper/rootvg-rootlv` will increase the size of the logical volume and the root filesystem by 10 GB and free up space. This command will help the administrator solve the problem and free up space. The other options are incorrect because they either do not affect the root filesystem (`fdisk -1 /dev/sdb`, `pvccreate /dev/sdb`, `lsblk /dev/sda`, or `vgextend /dev/rootvg /dev/sdb`) or do not use the correct syntax (`fdisk -1 /dev/sdb` instead of `fdisk -l /dev/sdb` or `parted -l /dev/mapper/rootvg-rootlv` instead of `parted /dev/mapper/rootvg-rootlv print`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 318-319, 331-332.

NEW QUESTION 52

A systems administrator created a new directory with specific permissions. Given the following output:

```
# file: comptia
# owner: root
# group: root user: : rwx group :: r-x other: :---
default:user :: rwx default:group :: r-x default:group:wheel: rwx default:mask :: rwx default:other ::-
Which of the following permissions are enforced on /comptia?
```

- A. Members of the wheel group can read files in /comptia.
- B. Newly created files in /comptia will have the sticky bit set.
- C. Other users can create files in /comptia.
- D. Only root can create files in /comptia.

Answer: A

Explanation:

The output shows the file access control list (FACL) of the /comptia directory, which is an extension of the standard Linux permissions that allows more fine-grained control over file and directory access¹. The FACL consists of two parts: the access ACL and the default ACL. The access ACL applies to the current object, while the default ACL applies to the objects created within the directory².

The access ACL has three entries: user, group, and other. These are similar to the standard Linux permissions, but they can be specified for individual users or groups as well. The user entry shows that the owner of the directory (root) has read, write, and execute permissions (rwx). The group entry shows that the group owner of the directory (root) has read and execute permissions (r-x). The other entry shows that all other users have no permissions (—).

The default ACL has five entries: user, group, group:wheel, mask, and other. These are applied to any files or directories created within /comptia. The user entry shows that the owner of the new object will have read, write, and execute permissions (rwx). The group entry shows that the group owner of the new object will have read and execute permissions (r-x). The group:wheel entry shows that the members of the wheel group will have read, write, and execute permissions (rwx) on the new object. The mask entry shows that the maximum permissions allowed for any user or group are read, write, and execute (rwx). The other entry shows that all other users will have no permissions (—) on the new object. Therefore, based on the FACL output, members of the wheel group can read files in /comptia, as they have read permission on both the directory and any files within it. Option B is incorrect because the sticky bit is not set on /comptia or any files within it.

The sticky bit is a special permission that prevents users from deleting or renaming files that they do not own in a shared directory³. It is symbolized by a `t` character in the execute position of others. Option C is incorrect because other users cannot create files in /comptia, as they have no permissions on the directory or any files within it. Option D is incorrect because root is not the only user who can create files in /comptia. Any user who has write permission on the directory can create files within it, such as members of the wheel group.

NEW QUESTION 55

Users have been unable to save documents to /home/tmp/temp and have been receiving the following error:

Path not found

A junior technician checks the locations and sees that /home/tmp/tempa was accidentally created instead of /home/tmp/temp. Which of the following commands should the technician use to fix this issue?

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

Answer: B

Explanation:

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

NEW QUESTION 58

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

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

Given the following analysis:

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

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

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

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

Which of the following is causing the error message?

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

Answer: D

Explanation:

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

NEW QUESTION 59

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

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

Answer: B

Explanation:

The `chgrp` command is used to change the group ownership of files and directories. By using this command, the administrator can assign the files in the newly created directory to the human resources group, which will allow the users in that group to access them. The other commands are not relevant for this task. For example:

? `chattr` is used to change the file attributes, such as making them immutable or append-only¹.

? `chage` is used to change the password expiration information for a user account².

? `chcon` is used to change the security context of files and directories, which is related to SELinux³.

References:

? The CompTIA Linux+ Certification Exam Objectives mention that the candidate should be able to “manage file and directory ownership and permissions” as part of the Hardware and System Configuration domain⁴.

? The web search result 2 explains how to use the `chgrp` command with examples.

? The web search result 3 compares the `chmod` and `chgrp` commands and their effects on file permissions.

NEW QUESTION 60

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 63

Which of the following tools is commonly used for creating CI/CD pipelines?

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

Answer: C

Explanation:

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

NEW QUESTION 68

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

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

Answer: A

Explanation:

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

NEW QUESTION 72

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 76

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

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

Answer: B

Explanation:

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

NEW QUESTION 80

A systems administrator is receiving tickets from users who cannot reach the application app that should be listening on port 9443/tcp on a Linux server. To troubleshoot the issue, the systems administrator runs netstat and receives the following output:

```
# netstat -anp | grep appd | grep -w LISTEN
tcp 0 0 127.0.0.1:9443 0.0.0.0:* LISTEN 1234/appd
```

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

- A. The IP address 0.0.0.0 is not valid.
- B. The application is listening on the loopback interface.
- C. The application is listening on port 1234.
- D. The application is not running.

Answer: B

Explanation:

The server is in a "Listen" state on port 9443 using its loopback address. The "1234" is a process-id. The cause of the issue is that the application is listening on the loopback interface. The loopback interface is a virtual network interface that is used for internal communication within the system. The loopback interface has the IP address 127.0.0.1, which is also known as localhost. The netstat output shows that the application is listening on port 9443 using the IP address 127.0.0.1. This means that the application can only accept connections from the same system, not from other systems on the network. This can prevent the users from reaching the application and cause the issue. The administrator should configure the application to listen on the IP address 0.0.0.0, which means all available interfaces, or on the specific IP address of the system that is reachable from the network. This will allow the application to accept connections from other systems and resolve the issue. The cause of the issue is that the application is listening on the loopback interface. This is the correct answer to the question. The other options are incorrect because they are not supported by the outputs. The IP address 0.0.0.0 is valid and means all interfaces, the application is not listening on port 1234, and the application is running as shown by the process ID 1234. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 383.

NEW QUESTION 83

A development team asks an engineer to guarantee the persistency of journal log files across system reboots. Which of the following commands would accomplish this task?

- A. `grep -i auto /etc/systemd/journald.conf && systemctl restart systemd-journald.service`
- B. `cat /etc/systemd/journald.conf | awk '(print $1,$3)'`
- C. `sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i 'persistent/s/^#/q' /etc/systemd/journald.conf`
- D. `journalctl --list-boots && systemctl restart systemd-journald.service`

Answer: C

Explanation:

The command `sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i 'persistent/s/^#/q' /etc/systemd/journald.conf` will accomplish the task of guaranteeing the persistency of journal log files across system reboots. The sed command is a tool for editing text files on Linux systems. The -i option modifies the file in place. The s command substitutes one string for another. The g flag replaces all occurrences of the string. The && operator executes the second command only if the first command succeeds. The q command quits after the first match. The /etc/systemd/journald.conf file is a configuration file for the systemd-journald service, which is responsible for collecting and storing log messages. The command `sed -i 's/auto/persistent/g' /etc/systemd/journald.conf` will replace the word auto with the word persistent in the file. This will change the value of the Storage option, which controls where the journal log files are stored. The value auto means that the journal log files are stored in the volatile memory and are lost after reboot, while the value persistent means that the journal log files are stored in the persistent storage and are preserved across reboots. The command `sed -i 'persistent/s/^#/q' /etc/systemd/journald.conf` will remove the # character at the beginning of the line that contains the word persistent. This will uncomment the Storage option and enable it. The command `sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i 'persistent/s/^#/q' /etc/systemd/journald.conf` will guarantee the persistency of journal log files across system reboots by changing and enabling the Storage option to persistent. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not change the value of the Storage option (`grep -i auto /etc/systemd/journald.conf && systemctl restart systemd-journald.service` or `cat /etc/systemd/journald.conf | awk '(print $1,$3)'`) or do not enable the Storage option (`journalctl --list-boots && systemctl restart systemd-journald.service`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Managing Logging and Monitoring, page 489.

NEW QUESTION 86

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

NEW QUESTION 89

A. ss -pint
B. tcpdump -nL
C. netstat -pn
D. lsof -lt

Answer: A

Explanation:

NEW QUESTION 92

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

Answer: B

Explanation:

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

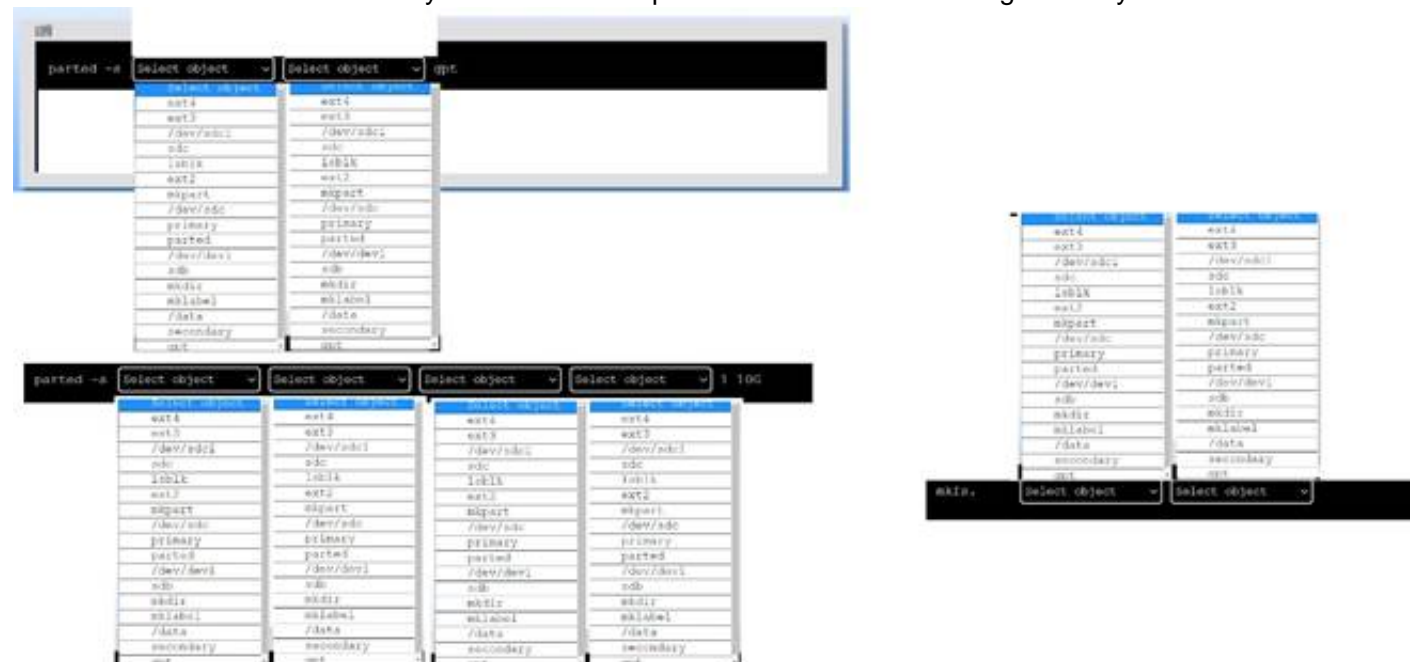
? D. `ssh accounts.pdf user@10.10.10.80:/tmp` will try to use the `ssh` command to log into the remote server, but it has the wrong syntax and arguments. The username should come before the remote host, and a file name is not a valid argument for `ssh`.

NEW QUESTION 97

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 mklable command, and the label type (gpt). The command is:

```
parted -s /dev/sdc mklable 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 101

An administrator attempts to connect to a remote server by running the following command:

```
$ nmap 192.168.10.36
```

Starting Nmap 7.60 (<https://nmap.org>) at 2022-03-29 20:20 UTC Nmap scan report for www1 (192.168.10.36)

Host is up (0.000091s latency). Not shown: 979 closed ports PORT STATE SERVICE 21/tcp open ftp 22/tcp filtered ssh 631/tcp open ipp

Nmap done: 1 IP address (1 host up) scanned in 0.06 seconds

Which of the following can be said about the remote server?

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

Answer: A

Explanation:

This is because the port 22/tcp is shown as filtered by nmap, which means that nmap cannot determine whether the port is open or closed because a firewall or other device is blocking its probes. If the SSH server was not running on the remote server, the port would be shown as closed, which means that nmap received a TCP RST packet in response to its probe. If the remote SSH server was using SSH protocol version 1, the port would be shown as open, which means that nmap received a TCP SYN/ACK packet in response to its probe. If the SSH host key on the remote server had expired, the port would also be shown as open, but the SSH client would display a warning message about the host key verification failure. Therefore, the best explanation for the filtered state of the port 22/tcp is that a firewall is preventing nmap from reaching the SSH server.

You can find more information about nmap port states and how to interpret them in the following web search results:

? Nmap scan what does STATE=filtered mean?

? How to find ports marked as filtered by nmap

? Technical Tip: NMAP scan shows ports as filtered

NEW QUESTION 105

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

```
# df -h /ftpusers/
```

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

```
# df -i /ftpusers/
```

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

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

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

Answer: C

Explanation:

The cause of the issue based on the output above is C. The inodes is at full capacity and would affect file creation for users.

An inode is a data structure that stores information about a file or directory, such as its name, size, permissions, owner, timestamps, and location on the disk. Each file or directory has a unique inode number that identifies it. The number of inodes on a filesystem is fixed when the filesystem is created, and it determines how many files and directories can be created on that filesystem. If the inodes are exhausted, no new files or directories can be created, even if there is enough disk space available.

The output for the second command shows that the /ftpusers/ filesystem has 0% of inodes available, which means that all the inodes have been used up. This would prevent users from creating new files on the FTP server. The administrator should either delete some unused files or directories to free up some inodes, or resize the filesystem to increase the number of inodes.

The other options are incorrect because:

* A. The users do not have the correct permissions to create files on the FTP server.

This is not true, because the output for the first command shows that the /ftpusers/ filesystem has 26% of disk space available, which means that there is enough

space for users to create files. The permissions of the files and directories are not shown in the output, but they are not relevant to the issue of inode exhaustion.

* B. The ftpusers filesystem does not have enough space.
This is not true, because the output for the first command shows that the /ftpusers/ filesystem has 26% of disk space available, which means that there is enough space for users to create files. The issue is not related to disk space, but to inode capacity.

* D. ftpusers is mounted as read only.
This is not true, because the output for the first command does not show any indication that the /ftpusers/ filesystem is mounted as read only. If it was, it would have an (ro) flag next to the mounted on column. A read only filesystem would prevent users from creating or modifying files on the FTP server, but it would not affect the inode usage.

NEW QUESTION 108

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 110

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 115

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

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

Answer: D

Explanation:

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

NEW QUESTION 120

A systems administrator wants to check for running containers. Which of the following commands can be used to show this information?

- A. docker pull
- B. docker stats
- C. docker ps
- D. docker list

Answer: C

Explanation:

The command that can be used to check for running containers is docker ps. The docker ps command can list all the containers that are currently running on the

system. To show all the containers, including those that are stopped, the administrator can use `docker ps -a`
References:

? [CompTIA Linux+ Study Guide], Chapter 11: Working with Containers, Section: Managing Containers with Docker
? [Docker PS Command with Examples]

NEW QUESTION 123

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 124

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 127

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

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

Answer: A

Explanation:

This command will allow the systems administrator to change the password of another user account in the system. The `sudo` prefix will grant the administrator the necessary privileges to perform this action, and the `passwd` command will prompt for the new password for the specified user. For example, if the administrator wants to change the password of a user named tom, the command will look like this:

sudo passwd tom

The other options are incorrect because:

* B. sudo userdel

This command will delete a user account from the system, not change its credentials. The userdel command removes the user's entry from the /etc/passwd and /etc/shadow files, as well as deletes the user's home directory and mail spool. This is not what the request asked for.

* C. sudo chage

This command will change the password expiration and aging information for a user account, not its credentials. The chage command can be used to set or modify various parameters related to password aging, such as the minimum and maximum number of days between password changes, the number of days before password expiration to issue a warning, and so on. This is not what the request asked for.

* D. sudo usermod

This command will modify various attributes of a user account, such as its login name, home directory, default shell, primary group, and so on. However, it cannot change the user's password directly. To do that, the usermod command requires the -p option followed by an encrypted password string, which is not easy to generate manually. Therefore, this is not a practical way to change a user's credentials.

References:

? How to Change Account Passwords on Linux

? How to Change a Password in Linux for Root and Other Users

? CompTIA Linux+ Certification Exam Objectives

NEW QUESTION 132

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

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

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

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

A. Reboot the serve

B. The volume will automatically go back to linear mode.

C. Replace the failed drive and reconfigure the mirror.

D. Reboot the serve

E. The volume will revert to stripe mode.

F. Recreate the logical volume.

Answer: B

Explanation:

The administrator should replace the failed drive and reconfigure the mirror to recover the volume. The LVM (Logical Volume Manager) is a tool for managing disk space on Linux systems. The LVM allows the administrator to create logical volumes that span across multiple physical volumes, such as hard disks or partitions. The LVM also supports different types of logical volumes, such as linear, striped, or mirrored. A mirrored logical volume is a type of logical volume that creates a copy of the data on another physical volume, providing redundancy and fault tolerance. The output shows that the logical volume is mirrored and that one leg of the mirror has failed due to the underlying physical volume. This means that one of the physical volumes that contains the data of the logical volume is damaged or missing. This can cause data loss and performance degradation. The administrator should replace the failed drive and reconfigure the mirror to recover the volume. The administrator should identify the failed physical volume by using commands such as pvdisplay, vgdisplay, or lvdisplay. The administrator should then remove the failed physical volume from the volume group by using the vgreduce command.

The administrator should then install a new drive and create a new physical volume by using the pvcreate command. The administrator should then add the new physical volume to the volume group by using the vgextend command. The administrator should then reconfigure the mirror by using the lvconvert command. The administrator should replace the failed drive and reconfigure the mirror to recover the volume. This is the correct answer to the question. The other options are incorrect because they either do not recover the volume (reboot the server. The volume will automatically go back to linear mode or reboot the server. The volume will revert to stripe mode) or do not preserve the data of the volume (recreate the logical volume). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 333-334.

NEW QUESTION 137

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

A. unzip -v

B. bzip2 -z

C. gzip

D. funzip

Answer: C

Explanation:

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

NEW QUESTION 140

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

A. sudo useradd -e 2021-09-30 Project_user

B. sudo useradd -c 2021-09-30 Project_user

C. sudo modinfo -F 2021-09-30 Project_uses

D. sudo useradd -m -d 2021-09-30 Project_user

Answer: A

Explanation:

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

NEW QUESTION 141

An administrator needs to get network information from a group of statically assigned workstations before they are reconnected to the network. Which of the following should the administrator use to obtain this information?

- A. `ip show`
- B. `ifcfg --a`
- C. `ifcfg --s`
- D. `i fname --s`

Answer: B

Explanation:

The `ifcfg` command is used to configure network interfaces on Linux systems. The `-a` option displays information about all network interfaces, including their IP addresses, netmasks, gateways, and other parameters. This command can help the administrator obtain the network information from the statically assigned workstations before they are reconnected to the network. References: [Linux Networking: `ifcfg` Command With Examples]

NEW QUESTION 143

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

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

Which of the following commands would resolve the issue?

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

Answer: D

Explanation:

The command that would resolve the issue is `chmod 600 .ssh/authorized_keys`. This command will change the permissions of the `.ssh/authorized_keys` file to 600, which means that only the owner of the file can read and write it. This is necessary for SSH key authentication to work properly, as SSH will refuse to use a key file that is accessible by other users or groups for security reasons. The output of `ls -l` shows that currently the `.ssh/authorized_keys` file has permissions of 664, which means that both the owner and group can read and write it, and others can read it.

The other options are not correct commands for resolving the issue. The `restorecon .ssh/authorized_keys` command will restore the default SELinux security context for the `.ssh/authorized_keys` file, but this will not change its permissions or ownership. The `ssh_keygen -t rsa -o .ssh/authorized_keys` command is invalid because `ssh_keygen` is not a valid command (the correct command is `ssh-keygen`), and the `-o` option is used to specify a new output format for the key file, not the output file name. The `chown root:root`

`.ssh/authorized_keys` command will change the owner and group of the `.ssh/authorized_keys` file to root, but this will not change its permissions or make it accessible by the user who wants to log in with SSH key authentication. References: How to Use Public Key Authentication with SSH; `chmod(1)` - Linux manual page

NEW QUESTION 144

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

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

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

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

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

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

Answer: B

Explanation:

The best task to perform to resolve this issue is B. Remove the entry for portal.comptia.org from the local hosts file. This is because the local hosts file has a wrong entry that maps portal.comptia.org to 10.10.10.55, which is different from the actual IP address of 192.168.1.55 that is returned by the DNS server. This causes a mismatch and prevents the workstation from accessing the website. By removing or correcting the entry in the hosts file, the workstation will use the DNS server to resolve the domain name and access the website successfully.

To remove or edit the entry in the hosts file, you need to have root privileges and use a text editor such as vi or nano. For example, you can run the command:

```
sudo vi /etc/hosts
```

and delete or modify the line that says: 10.10.10.55 portal.comptia.org

Then save and exit the file.

NEW QUESTION 149

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 151

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

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

Answer: D

NEW QUESTION 154

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

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

Answer: C

Explanation:

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

NEW QUESTION 157

A systems administrator needs to reconfigure a Linux server to allow persistent IPv4 packet forwarding. Which of the following commands is the correct way to accomplish this task?

- A. echo 1 > /proc/sys/net/ipv4/ipv_forward
- B. sysctl -w net.ipv4.ip_forward=1
- C. firewall-cmd --enable ipv4_forwarding
- D. systemctl start ipv4_forwarding

Answer: B

Explanation:

The command sysctl -w net.ipv4.ip_forward=1 enables IPv4 packet forwarding temporarily by setting the kernel parameter net.ipv4.ip_forward to 1. To make this change persistent, the administrator needs to edit the file /etc/sysctl.conf and add the line net.ipv4.ip_forward = 1. The other options are incorrect because they either use the wrong file (/proc/sys/net/ipv4/ipv_forward), the wrong command (firewall- cmd or systemctl), or the wrong option (--enable or start). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 378.

NEW QUESTION 161

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

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

Answer: D

Explanation:

The command `ip route add 10.0.213.5/32 via 10.0.5.1 dev eth0` adds a static route to the routing table that sends packets destined for 10.0.213.5/32 (a single host) through the gateway 10.0.5.1 on the interface eth0. This is the correct way to achieve the goal. The other options are incorrect because they either use the wrong syntax (`route -i eth0 -p add`), the wrong command (`route modify`), or the wrong file (`/proc/net/route`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 379.

NEW QUESTION 163

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 165

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 169

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

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

Answer: B

Explanation:

The command `iptables -t filter -A INPUT -p tcp --dport 4000:5000 -j ACCEPT` will enforce the rule of allowing incoming traffic to ports in the range of 4000 to 5000 on a Linux server. The `iptables` command is a tool for managing firewall rules on Linux systems. The `-t` option specifies the table to operate on, in this case `filter`, which is the default table that contains the rules for filtering packets. The `-A` option appends a new rule to the end of a chain, in this case `INPUT`, which is the chain that processes the packets that are destined for the local system. The `-p` option specifies the protocol to match, in this case `tcp`, which is the transmission control protocol. The `--dport` option specifies the destination port or port range to match, in this case `4000:5000`, which is the range of ports from 4000 to 5000. The `-j` option specifies the target to jump to if the rule matches, in this case `ACCEPT`, which is the target that allows the packet to pass through.

The command `iptables -t filter -A INPUT -p tcp --dport 4000:5000 -j ACCEPT` will add a new rule to the end of the `INPUT` chain that will accept the incoming TCP packets that have a destination port between 4000 and 5000. This command will enforce the rule and allow the traffic to the specified ports. This is the correct command to use to accomplish the task. The other options are incorrect because they either use the wrong options (`-f` instead of `-t` or `-D` instead of `-A`) or do not exist (`iptables filter -A INPUT -p tcp --dport 4000:5000 -D ACCEPT` or `iptables filter -S INPUT -p tcp --dport 4000:5000 -A ACCEPT`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 543.

NEW QUESTION 173

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 178

A Linux system is failing to boot. The following error is displayed in the serial console: `[[1;33mDEPEND[Om] Dependency failed for /data.`

`[[1;33mDEPEND[Om] Dependency failed for Local File Systems`

...

Welcome to emergency mode! After logging in, type `"journalctl -xb"` to view system logs,

`"systemctl reboot"` to reboot, `"systemctl default"` to try again to boot into default mode.

Give root password for maintenance (or type `Control-D` to continue)

Which of the following files will need to be modified for this server to be able to boot again?

- A. `/etc/mtab`
- B. `/dev/sda`
- C. `/etc/fstab`
- D. `/etc/grub.conf`

Answer: C

Explanation:

The file that will need to be modified for the server to be able to boot again is `/etc/fstab`. The `/etc/fstab` file is a file that contains the information about the file systems that are mounted at boot time on Linux systems. The file specifies the device name, mount point, file system type, mount options, dump frequency, and pass number for each file system. The error message indicates that the dependency failed for `/data`, which is a mount point for a file system. This means that the system could not mount the `/data` file system at boot time, which caused the system to enter the emergency mode. The emergency mode is a mode that allows the administrator to log in as the root user and perform basic tasks such as repairing the system. The administrator should modify the `/etc/fstab` file and check the entry for the `/data` file system. The administrator should look for any errors or inconsistencies in the device name, file system type, or mount options, and correct them. The administrator should also verify that the device and the file system are intact and functional by using commands such as `blkid`, `fdisk`, `fsck`, or `mount`. The administrator should then reboot the system and see if the issue is resolved. The file that will need to be modified for the server to be able to boot again is `/etc/fstab`. This is the correct answer to the question. The other options are incorrect because they are not related to the file systems that are mounted at boot time (`/etc/mtab`, `/dev/sda`,

or `/etc/grub.conf`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 321.

NEW QUESTION 179

A Linux administrator is configuring a two-node cluster and needs to be able to connect the nodes to each other using SSH keys from the root account. Which of

the following commands will accomplish this task?

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

Answer: C

Explanation:

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

NEW QUESTION 181

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

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

Answer: C

Explanation:

The mechanism that enables administrators to configure and enforce MFA on a Linux system is PAM. PAM stands for Pluggable Authentication Modules, which is a framework for managing authentication and authorization on Linux systems. PAM allows the administrator to define the rules and policies for accessing various system resources and services, such as login, sudo, ssh, or cron. PAM also supports different types of authentication methods, such as passwords, tokens, biometrics, or smart cards. PAM can be used to implement MFA, which stands for Multi-Factor Authentication, which is a security technique that requires the user to provide more than one piece of evidence to prove their identity. MFA can enhance the security of the system and prevent unauthorized access. PAM enables administrators to configure and enforce MFA on a Linux system. This is the correct answer to the question. The other options are incorrect because they either do not manage authentication and authorization on Linux systems (Kerberos or PKI) or do not support MFA (SELinux). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Implementing Basic Security, page 517.

NEW QUESTION 183

Which of the following would significantly help to reduce data loss if more than one drive fails at the same time?

- A. Server clustering
- B. Load balancing
- C. RAID
- D. VDI

Answer: C

Explanation:

RAID stands for Redundant Array of Independent Disks, which is a technology that combines multiple physical disks into a logical unit that provides improved performance, reliability, or both. RAID can significantly help to reduce data loss if more than one drive fails at the same time, depending on the RAID level used. For example, RAID 1 (mirroring) duplicates the data on two or more disks, so that if one disk fails, the data can be recovered from another disk. RAID 5 (striping with parity) distributes the data and parity information across three or more disks, so that if one disk fails, the data can be reconstructed from the remaining disks. RAID 6 (striping with double parity) extends RAID 5 by adding another parity block, so that if two disks fail, the data can still be recovered from the remaining disks. References: [What is RAID?]

NEW QUESTION 186

A systems administrator is checking the system logs. The administrator wants to look at the last 20 lines of a log. Which of the following will execute the command?

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

Answer: B

Explanation:

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

NEW QUESTION 188

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

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

D. apt remove packagename

Answer: A

Explanation:

The command that can be used to remove an RPM package that was installed by mistake is `dnf remove packagename`. This command will use the DNF package manager to uninstall an RPM package and its dependencies from a Linux system that uses RPM-based distributions, such as Red Hat Enterprise Linux or CentOS. The DNF package manager handles dependency resolution and metadata searching for RPM packages.

The other options are not correct commands for removing an RPM package from a Linux system. The `apt-get remove packagename` and `apt remove packagename` commands are used to remove Debian packages from a Linux system that uses Debian-based distributions, such as Ubuntu or Debian. They are not compatible with RPM packages. The `rpm -i packagename` command is used to install an RPM package, not to remove it. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 9: Managing Software Packages; How to install/remove/query/update RPM packages in Linux (Cheat Sheet ...

NEW QUESTION 190

The applications team is reporting issues when trying to access the web service hosted in a Linux system. The Linux systems administrator is reviewing the following outputs:

Output 1:

* httpd.service = The Apache HTTPD Server

Loaded: loaded (/usr/lib/systemd/system/httpd.service; disabled; vendor preset: disabled) Active: inactive (dead)

Docs: man:httpd(8) man:apachectl(8) Output 2:

16:51:16 up 28 min, 1 user, load average: 0.00, 0.00, 0.07

Which of the following statements best describe the root cause? (Select two).

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

Answer: CD

Explanation:

The `httpd.service` is the Apache HTTPD Server, which is a web service that runs on Linux systems. The output 1 shows that the `httpd.service` is inactive (dead), which means that it is not running. The output 1 also shows that the `httpd.service` is disabled, which means that it is not enabled to auto start at boot time.

Therefore, the statements C and D best describe the root cause of the issue. The statements A, B, E, and F are incorrect because they do not match the output 1.

References: [How to Manage Systemd Services on a Linux System]

NEW QUESTION 191

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

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

Answer: C

Explanation:

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

NEW QUESTION 194

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

- A. `systemctl cancel nginx`
- B. `systemctl disable nginx`
- C. `systemctl mask nginx`
- D. `systemctl stop nginx`

Answer: C

Explanation:

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

NEW QUESTION 196

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 198

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

- A. `find /etc/passwd —size +500`
- B. `cut —d: f1 / etc/ passwd > 500`
- C. `awk -F: '$3 > 500 {print $1}' /etc/passwd`
- D. `sed '/UID/' /etc/passwd < 500`

Answer: C

Explanation:

The correct command to list all local accounts in which the UID is greater than 500 is:

`awk -F: '$3 > 500 {print $1}' /etc/passwd`

This command uses awk to process the `/etc/passwd` file, which contains information about the local users on the system. The `-F:` option specifies that the fields are separated by colons. The `$3` refers to the third field, which is the UID. The condition `$3 > 500` filters out the users whose UID is greater than 500. The action `{print $1}` prints the first field, which is the username.

The other commands are incorrect because:

? `find /etc/passwd —size +500` will search for files that are larger than 500 blocks in size, not users with UID greater than 500.

? `cut —d: f1 / etc/ passwd > 500` will cut the first field of the `/etc/passwd` file using colon as the delimiter, but it will not filter by UID or print only the usernames. The `> 500` part will redirect the output to a file named 500, not compare with the UID.

? `sed '/UID/' /etc/passwd < 500` will use sed to edit the `/etc/passwd` file and replace any line that contains UID with 500, not list the users with UID greater than 500.

The `< 500` part will redirect the input from a file named 500, not compare with the UID.

References:

? `Linux List All Users In The System Command` - nixCraft, section “List all users in Linux using `/etc/passwd` file”.

? `Unix script getting users with UID bigger than 500` - Stack Overflow, section “Using awk”.

NEW QUESTION 202

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

- A. `systemctl stop sshd`
- B. `systemctl mask sshd`
- C. `systemctl reload sshd`
- D. `systemctl start sshd`

Answer: C

Explanation:

The `systemctl reload sshd` command can be used to apply the configuration changes of the SSH server daemon without restarting it. This is useful to avoid interrupting existing connections. The `systemctl stop sshd` command would stop the SSH server daemon, not apply the changes. The `systemctl mask sshd` command would prevent the SSH server daemon from being started, not apply the changes. The `systemctl start sshd` command would start the SSH server

daemon if it is not running, but it would not apply the changes if it is already running. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Secure Shell (SSH), page 415.

NEW QUESTION 205

Several users reported that they were unable to write data to the /oracle1 directory. The following output has been provided:

Filesystem	Size	Used	Available	Use%	Mounted on
/dev/sdb1	100G	50G	50G	50%	/oracle1

Which of the following commands should the administrator use to diagnose the issue?

- A. df -i /oracle1
- B. fdisk -l /dev/sdb1
- C. lsblk /dev/sdb1
- D. du -sh /oracle1

Answer: A

Explanation:

The administrator should use the command `df -i /oracle1` to diagnose the issue of users being unable to write data to the /oracle1 directory. This command will show the inode usage of the /oracle1 filesystem, which indicates how many files and directories can be created on it. If the inode usage is 100%, it means that no more files or directories can be added, even if there is still free space on the disk. The administrator can then delete some unnecessary files or directories, or increase the inode limit of the filesystem, to resolve the issue.

The other options are not correct commands for diagnosing this issue. The `fdisk -l /dev/sdb1` command will show the partition table of /dev/sdb1, which is not relevant to the inode usage. The `lsblk /dev/sdb1` command will show information about /dev/sdb1 as a block device, such as its size, mount point, and type, but not its inode usage. The `du -sh /oracle1` command will show the disk usage of /oracle1 in human-readable format, but not its inode usage. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 7: Managing Disk Storage; How to Check Inode Usage in Linux - Fedingo

NEW QUESTION 209

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

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

Answer: C

Explanation:

The docker command is used to manage Docker containers, images, networks, volumes, and other resources on a Linux system. Docker is a platform that allows users to run applications in isolated environments called containers. Docker also provides networking features that allow users to create and manage networks for containers.

To delete all unused networks that are not referenced by any container, the cloud engineer can use the `docker network prune` command. This command will remove all networks that have no containers connected to them. The statement C is correct.

The statements A, B, and D are incorrect because they do not delete all unused networks.

The `docker network erase` and `docker network clear` commands do not exist. The `docker network rm` command deletes a specific network by name or ID, but not all unused networks. References: [How to Manage Docker Networks]

NEW QUESTION 213

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

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

Answer: B

Explanation:

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

NEW QUESTION 216

A systems administrator is troubleshooting a connectivity issue pertaining to access to a system named db.example.com. The system IP address should be 192.168.20.88. The administrator issues the dig command and receives the following output:

```
;; ANSWER SECTION:
db.example.com.      15 IN A 192.168.20.89
```

The administrator runs `grep db.example.com /etc/hosts` and receives the following output:

```
192.168.20.89 db.example.com
```

Given this scenario, which of the following should the administrator do to address this issue?

- A. Modify the /etc/hosts file and change the db.example.com entry to 192.168.20.89.
- B. Modify the /etc/network file and change the db.example.com entry to 192.168.20.88.
- C. Modify the /etc/network file and change the db.example.com entry to 192.168.20.89.
- D. Modify the /etc/hosts file and change the db.example.com entry to 192.168.20.88.

Answer: D

Explanation:

The administrator should modify the /etc/hosts file and change the db.example.com entry to 192.168.20.88 to address the issue. The /etc/hosts file is a file that maps hostnames to IP addresses on Linux systems. The file can be used to override the DNS resolution and provide a local lookup for hostnames. The dig output shows that the DNS returns the IP address 192.168.20.88 for the hostname db.example.com, which is the correct IP address of the system. The grep output shows that the /etc/hosts file contains an entry for db.example.com with the IP address 192.168.20.89, which is the wrong IP address of the system. This can cause a conflict and prevent the system from being accessed by the hostname. The administrator should modify the /etc/hosts file and change the db.example.com entry to 192.168.20.88, which is the correct IP address of the system. This will align the /etc/hosts file with the DNS and allow the system to be accessed by the hostname. The administrator should modify the /etc/hosts file and change the db.example.com entry to 192.168.20.88 to address the issue. This is the correct answer to the question. The other options are incorrect because they either do not modify the /etc/hosts file (modify the /etc/network file and change the db.example.com entry to 192.168.20.88 or modify the /etc/network file and change the db.example.com entry to 192.168.20.89) or do not change the IP address to the correct one (modify the /etc/hosts file and change the db.example.com entry to 192.168.20.89). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 378.

NEW QUESTION 221

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 225

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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