

CompTIA

Exam Questions XK0-005

CompTIA Linux+ Certification Exam



NEW QUESTION 1

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

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

Answer: D

Explanation:

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

NEW QUESTION 2

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

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

Answer: B

Explanation:

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

NEW QUESTION 3

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

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 4

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 6

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 7

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

Hostname: devel.comptia.org

IP address: 5.5.5.1, 5.5.5.2, 5.5.5.3, 5.5.5.4

Name server: 5.5.5.254

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

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

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

Answer: BDE

Explanation:

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

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

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

? NS: This record type is used to delegate a domain or a subdomain to another name server. The administrator needs one NS record for comptia.org, pointing to 5.5.5.254, which is the name server that hosts the records for the subdomain devel.comptia.org². This will allow users to resolve the hostnames under comptia.org by querying the name server 5.5.5.254².

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

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

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

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

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

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

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

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

NEW QUESTION 8

A cloud engineer needs to block the IP address 192.168.10.50 from accessing a Linux server. Which of the following commands will achieve this goal?

- A. iptables -F INPUT -j 192.168.10.50 -m DROP
- B. iptables -A INPUT -s 192.168.10.50 -j DROP
- C. iptables -i INPUT --ipv4 192.168.10.50 -z DROP
- D. iptables -j INPUT 192.168.10.50 -p DROP

Answer: B

Explanation:

The correct command to block the IP address 192.168.10.50 from accessing a Linux server is iptables -A INPUT -s 192.168.10.50 -j DROP. This command appends a rule to the INPUT chain that matches the source address 192.168.10.50 and jumps to the DROP target, which discards the packet. The other commands are incorrect because they either have invalid syntax, wrong parameters, or wrong order of arguments. References: CompTIA Linux+ Study Guide, Fourth Edition, page 457-458.

NEW QUESTION 9

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

- A. passwd
- B. ssh
- C. ssh-keygen
- D. pwgen

Answer: C

Explanation:

The command that would allow a user to choose a stronger password and set it on the existing SSH key file is ssh-keygen -p -f <keyfile>. This command uses the ssh-keygen tool, which is used to generate, manage, and convert authentication keys for SSH. The -p option stands for passphrase, and it allows the user to change or remove the passphrase of an existing private key file. The -f option specifies the filename of the key file. The command will prompt the user for the old passphrase, and then for the new passphrase twice.

The other options are not correct commands for changing the password of an SSH key file. The passwd command is used to change the password of a user account on a Linux system, not an SSH key file. The ssh command is used to log in to a remote system using SSH, not to change the password of an SSH key file. The pwgen command is used to generate random passwords, not to change the password of an SSH key file.

References: ssh-keygen(1) - Linux manual page; How To: Change Passphrase for SSH Private Key - Unix Tutorial

NEW QUESTION 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 session?

- 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 10

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


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

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

- A.
- ```
Iptables -A FORWARD -i eth0 -p tcp --dport 80 -j ACCEPT
Iptables -A FORWARD -i eth0 -p tcp --dport 22 -j ACCEPT
```
- B.
- ```
Iptables -A INPUT -i eth0 -p tcp --dport 80 -j ACCEPT
Iptables -A INPUT -i eth0 -p tcp --dport 22 -j ACCEPT
```
- C.
- ```
Iptables -A INPUT -i eth0 -p tcp --sport 80 -j ACCEPT
Iptables -A INPUT -i eth0 -p tcp --sport 22 -j ACCEPT
```
- D.
- ```
Iptables -A INPUT -i eth0 -p tcp --dport :80 -j ACCEPT
Iptables -A INPUT -i eth0 -p tcp --dport :22 -j ACCEPT
```

Answer: A

Explanation:

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

NEW QUESTION 12

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 13

A systems administrator needs to clone the partition /dev/sdc1 to /dev/sdd1. Which of the following commands will accomplish this task?

- A. tar -cvzf /dev/sdd1 /dev/sdc1
- B. rsync /dev/sdc1 /dev/sdd1
- C. dd if=/dev/sdc1 of=/dev/sdd1
- D. scp /dev/sdc1 /dev/sdd1

Answer: C

Explanation:

The command dd if=/dev/sdc1 of=/dev/sdd1 copies the data from the input file (if) /dev/sdc1 to the output file (of) /dev/sdd1, byte by byte. This is the correct way to clone a partition. The other options are incorrect because they either compress the data (tar -cvzf), synchronize the files (rsync), or copy the files over a network (scp), which are not the same as cloning a partition. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 321.

NEW QUESTION 14

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 16

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

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

Answer: A

Explanation:

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

? C. systemctl stop httpd will only stop the service if it is currently running, but it will not prevent it from being started again.

? D. systemctl reload httpd will only reload the configuration files of the service, but it will not stop or disable it.

NEW QUESTION 21

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 23

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 26

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

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

Answer: A

Explanation:

The correct answer is A. chmod g+s /home/test

This command will set the setgid bit on the /home/test directory, which means that any file or subdirectory created in the directory will inherit the group ownership of the directory. This way, the group permissions on files created in the directory will be preserved. The chmod command is used to change the permissions of files and directories. The g+s option is used to set the setgid bit for the group.

The other options are incorrect because:

* B. chgrp test /home/test

This command will change the group ownership of the /home/test directory to test, but it will not affect the group ownership of files created in the directory. The chgrp command is used to change the group of files and directories. The test /home/test arguments are used to specify the new group and the target directory.

* C. chmod 777 /home/test

This command will give read, write, and execute permissions to everyone (owner, group, and others) on the /home/test directory, but it will not affect the group ownership or permissions of files created in the directory. The chmod command is used to change the permissions of files and directories. The 777 argument is an octal number that represents the permissions in binary form.

* D. chown -hR test /home/test

This command will change the owner and group of the /home/test directory and all its contents recursively to test, but it will not preserve the original group permissions on files created in the directory. The chown command is used to change the owner and group of files and directories. The -hR option is used to affect symbolic links and operate on all files and directories recursively. The test /home/test arguments are used to specify the new owner and group and the target directory.

References:

? How to Set File Permissions Using chmod

? How to Use Chmod Command in Linux with Examples

? How to Use Chown Command in Linux with Examples

? [How to Use Chgrp Command in Linux with Examples]

NEW QUESTION 31

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

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

Answer: D

Explanation:

To allow the same Kubernetes container configurations to be deployed in different environments without statically configuring custom locations, the engineer can use an ambassador container (D). An ambassador container is a proxy container that handles communication between containers and external services. It can dynamically configure locations based on environment variables or other methods. The other options are not related to this requirement. References:

? [CompTIA Linux+ Study Guide], Chapter 11: Working with Containers, Section: Using Ambassador Containers

? [How to Use Ambassador Containers]

NEW QUESTION 32

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

- A. git clone <https://git.company.com/admin/project.git>

- B. git checkout <https://git.company.com/admin/project.git>
- C. git pull <https://git.company.com/admin/project.git>
- D. git branch <https://git.company.com/admin/project.git>

Answer: A

Explanation:

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

NEW QUESTION 33

Users report that connections to a MariaDB service are being closed unexpectedly. A systems administrator troubleshoots the issue and finds the following message in /var/log/messages:

```
dbserver kernel: out of Memory: Killed process 1234 (mysqld).
```

Which of the following is causing the connection issue?

- A. The process mysqld is using too many semaphores.
- B. The server is running out of file descriptors.
- C. Something is starving the server resources.
- D. The amount of RAM allocated to the server is too high.

Answer: B

Explanation:

The message in /var/log/messages indicates that the server is running out of file descriptors. A file descriptor is a non-negative integer identifier for an open file in Linux. Each process has a table of open file descriptors where a new entry is appended upon opening a new file. There is a limit on how many file descriptors a process can open at a time, which depends on the system configuration and the user privileges. If a process tries to open more files than the limit, it will fail with an error message like "Too many open files". This could cause connections to be closed unexpectedly or other problems with the application. The other options are not correct causes for the connection issue. The process mysqld is not using too many semaphores, which are synchronization mechanisms for processes that share resources. Semaphores are not related to file descriptors or open files. Something is not starving the server resources, which could mean high CPU usage, memory pressure, disk I/O, network congestion, or other factors that affect performance. These could cause slowdowns or timeouts, but not file descriptor exhaustion. The amount of RAM allocated to the server is not too high, which could cause swapping or paging if it exceeds the physical memory available. This could also affect performance, but not file descriptor availability. References: File Descriptor Requirements (Linux Systems); Limits on the Number of Linux File Descriptors

NEW QUESTION 34

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 35

A Linux systems administrator is configuring a new filesystem that needs the capability to be mounted persistently across reboots. Which of the following commands will accomplish this task? (Choose two.)

- A. df -h /data
- B. mkfs.ext4 /dev/sdc1
- C. fsck /dev/sdc1
- D. fdisk -l /dev/sdc1
- E. echo "/data /dev/sdc1 ext4 defaults 0 0" >> /etc/fstab
- F. echo "/dev/sdc1 /data ext4 defaults 0 0" >> /etc/fstab

Answer: BF

Explanation:

"modify the /etc/fstab text file to automatically mount the new partition by opening it in an editor and adding the following line:
/dev/ xxx 1 /data ext4 defaults 1 2
where xxx is the device name of the storage device"

<https://learning.oreilly.com/library/view/mastering-linux-system/9781119794455/b01.xhtml> To configure a new filesystem that needs the capability to be mounted persistently across reboots, two commands are needed: `mkfs.ext4 /dev/sdc1` and `echo "/dev/sdc1 /data ext4 defaults 0 0" >> /etc/fstab`. The first command creates an ext4 filesystem on the device `/dev/sdc1`, which is the partition that will be used for the new filesystem. The second command appends a line to the `/etc/fstab` file, which is the configuration file that controls persistent mount points of filesystems. The line specifies the device name, the mount point (`/data`), the filesystem type (ext4), the mount options (defaults), and the dump and pass values (0 0). The other commands are incorrect because they either do not create or configure a filesystem, or they have wrong syntax or arguments. References: CompTIA Linux+ Study Guide, Fourth Edition, page 409-410, 414-415.

NEW QUESTION 38

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

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

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

- A. `gpg /dev/sdc1`
- 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/ sdc1`

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 41

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

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

Answer: D

Explanation:

The command that will allow the administrator to determine the first command that is executed inside the container right after it starts is `docker inspect <container_id>`. This command will display detailed information about the container, including its configuration, state, network settings, mounts, and logs. One of the configuration fields is “Entrypoint”, which shows the command that is executed when the container is run. The entrypoint can be specified in the Dockerfile or overridden at runtime using the `--entrypoint` option.

The other options are not correct commands for determining the first command that is executed inside the container. The `docker export <container_id>` command will export the contents of the container’s filesystem as a tar archive to STDOUT. This will not show the entrypoint of the container, but only its files. The `docker info <container_id>` command is invalid because `docker info` does not take any arguments. It shows system-wide information about Docker, such as the number of containers, images, volumes, networks, and storage drivers. The `docker start <container_id>` command will start a stopped container and attach its STDOUT and STDERR to the terminal. This will not show the entrypoint of the container, but only its output. References: `docker inspect` | Docker Docs; `docker export` | Docker Docs; `docker info` | Docker Docs; `docker start` | Docker Docs

NEW QUESTION 43

A junior administrator is setting up a new Linux server that is intended to be used as a router at a remote site. Which of the following parameters will accomplish this goal?

A.

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

A.

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

B.

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

C.

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

Answer: C

Explanation:

The parameter `net.ipv4.ip_forward=1` will accomplish the goal of setting up a new Linux server as a router. This parameter enables the IP forwarding feature, which allows the server to forward packets between different network interfaces. This is necessary for a router to route traffic between different networks. The parameter can be set in the `/etc/sysctl.conf` file or by using the `sysctl` command. This is the correct parameter to use to accomplish the goal. The other options are incorrect because they either do not exist (`net.ipv4.ip_forwarding` or `net.ipv4.ip_route`) or do not enable IP forwarding (`net.ipv4.ip_forward=0`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 382.

NEW QUESTION 44

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 47

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 50

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

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

Answer: A

Explanation:

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

NEW QUESTION 52

A systems administrator is compiling a report containing information about processes that are listening on the network ports of a Linux server. Which of the following commands will allow the administrator to obtain the needed information?

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

Answer: A

Explanation:

The command `ss -pint` will allow the administrator to obtain the needed information about processes that are listening on the network ports of a Linux server. The `ss` command is a tool for displaying socket statistics on Linux systems. Sockets are endpoints of network communication that allow processes to exchange data over the network. The `ss` command can show various information about the sockets, such as the state, address, port, protocol, and process. The `-pint` option specifies the filters and flags that the `ss` command should apply. The `-p` option shows the process name and ID that owns the socket. The `-i` option shows the internal information about the socket, such as the send and receive queue, the congestion window, and the retransmission timeout. The `-n` option shows the numerical address and port, instead of resolving the hostnames and service names. The `-t` option shows only the TCP sockets, which are the most common type of sockets used for network communication. The command `ss -pint` will display the socket statistics for the TCP sockets, along with the process name and ID, the numerical address and port, and the internal information. This will allow the administrator to obtain the needed information about processes that are listening on the network ports of a Linux server. This is the correct command to use to obtain the needed information. The other options are incorrect because they either do not show the socket statistics (`tcpdump -nL` or `lsof -lt`) or do not show the process name and ID (`netstat -pn`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 389.

NEW QUESTION 57

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

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

Answer: B

Explanation:

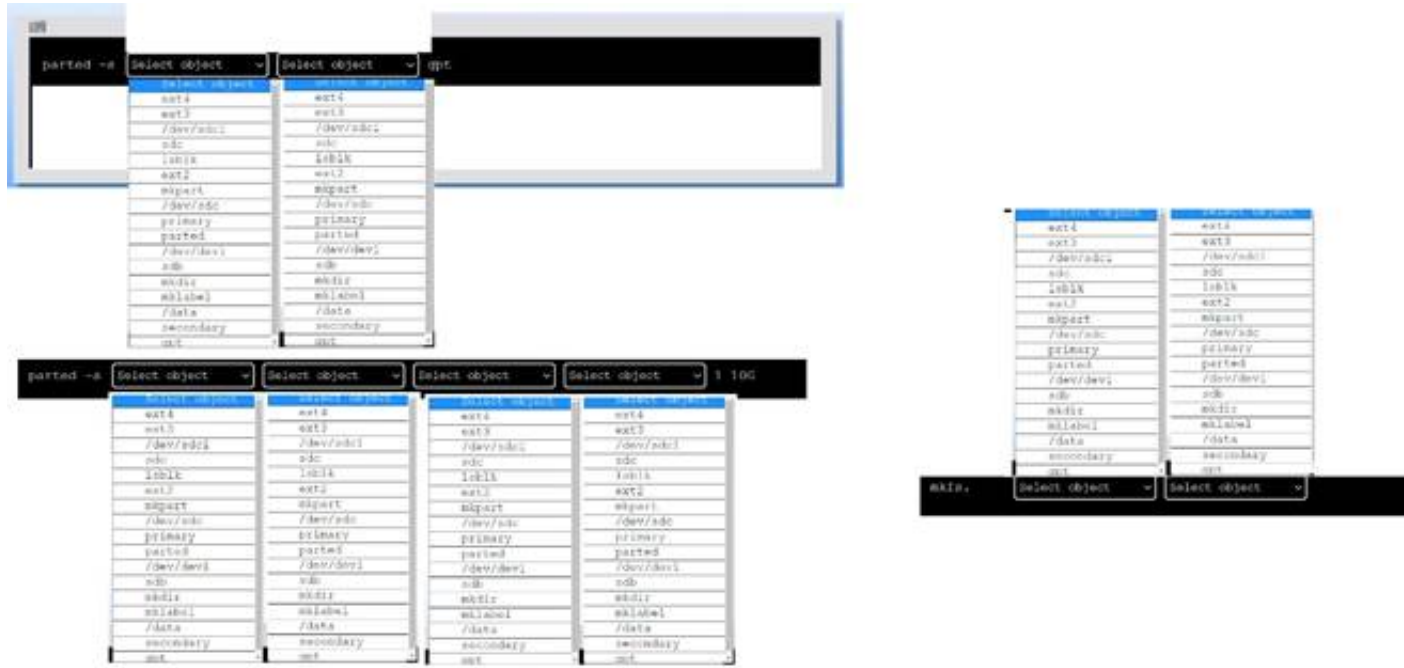
The best command to use to transfer the local file `accounts.pdf` to the remote /`tmp` directory of the server with the IP address `10.10.10.80` is B. `scp accounts.pdf user@10.10.10.80:/tmp`. This command will use the secure copy protocol (`scp`) to copy the file from the local machine to the remote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

NEW QUESTION 62

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 63

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

- A. nslookup
 B. rsyn
 C. netstat
 D. host

Answer: A

Explanation:

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

NEW QUESTION 66

What is the main objective when using Application Control?

- A. To filter out specific content.
 B. To assist the firewall blade with handling traffic.
 C. To see what users are doing.
 D. Ensure security and privacy of information.

Answer: D

Explanation:

The main objective when using Application Control is to ensure the security and privacy of information. Application Control is a security practice that blocks or restricts unauthorized applications from executing in ways that put data at risk. The control functions vary based on the business purpose of the specific application, but the main objective is to help ensure the privacy and security of data used by and transmitted between applications¹. Application Control can also prevent malware, untrusted, or unwanted applications from running on the network, reducing the risks and costs associated with data breaches¹. Application Control can also improve the overall network stability and performance by eliminating unnecessary or harmful applications¹. Application Control is not mainly used to filter out specific content, although it can be combined with other technologies such as URL filtering or content filtering to achieve that goal. Application Control is not mainly used to assist the firewall blade with handling traffic, although it can be integrated with firewall policies to enforce granular access rules based on applications. Application Control is not mainly used to see what users are doing, although it can provide visibility and

reporting on application usage and activity.

NEW QUESTION 71

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 76

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 77

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

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

Answer: D

Explanation:

The command that should be run on the DNS forwarder server to

accomplish the task is `firewall-cmd --zone=public --add-port=53/udp --permanent`.

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

NEW QUESTION 81

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 83

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

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

Answer: C

Explanation:

The command `sudo fdisk -l` should be issued to verify the device name of the partition. The `sudo` command allows the administrator to run commands as the superuser or another user. The `fdisk` command is a tool for manipulating disk partitions on Linux systems. The `-l` option lists the partitions on all disks or a specific disk. The command `sudo fdisk -l` will show the device names, sizes, types, and other information of the partitions on all disks. The administrator can identify the device name of the partition by looking at the output. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not list the partitions (`sudo fdisk /dev/sda` or `sudo fdisk -h`) or do not exist (`sudo fdisk -s /dev/sda`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 317.

NEW QUESTION 84

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 89

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 `pvdiskdisplay`, `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 `pvccreate` 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 91

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

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

Answer: D

Explanation:

The command `find . -type f -print | xargs grep -li denied` will accomplish the task of identifying files that contain any occurrence of the word denied. The `find` command is a tool for searching for files and directories on Linux systems. The `.` is the starting point of the search, which means the current directory. The `-type f` option specifies the type of the file, which means regular file. The `-print` option prints the full file name on the standard output. The `|` is a pipe symbol that redirects the output of one command to the input of another command. The `xargs` command is a tool for building and executing commands from standard input. The `grep` command is a tool for searching for patterns in files or input.

The `-li` option specifies the flags that the `grep` command should apply. The `-l` flag shows only the file names that match the pattern, instead of the matching lines. The `-i` flag ignores the case of the pattern, which means it matches both uppercase and lowercase letters.

The `denied` is the pattern that the `grep` command should search for. The command `find . -type f -print | xargs grep -li denied` will find all the regular files in the current directory and its subdirectories, and then search for any occurrence of the word denied in those files, ignoring the case, and print only the file names that match the pattern. This will allow the administrator to identify files that contain any occurrence of the word denied. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not ignore the case of the pattern (`find . -type f -print | xargs grep -ln denied` or `find . -type f -print | xargs grep -wL denied`) or do not show the file names that match the pattern (`find . -type f -print | xargs grep -nv denied`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Managing Logging and Monitoring, page 489.

NEW QUESTION 93

An administrator would like to list all current containers, regardless of their running state. Which of the following commands would allow the administrator to accomplish this task?

- A. `docker ps -a`
- B. `docker list`
- C. `docker image ls`
- D. `docker inspect image`

Answer: A

Explanation:

The best command to use to list all current containers, regardless of their running state, is A. `docker ps -a`. This command will show all containers, both running and stopped, with details such as container ID, image name, status, and ports. The other commands are either invalid or not relevant for this task. For example:

? B. `docker list` is not a valid command. There is no subcommand named `list` in `docker`.

? C. `docker image ls` will list all the images available on the local system, not the containers.

? D. `docker inspect image` will show detailed information about a specific image, not all the containers.

NEW QUESTION 98

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 99

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 102

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 104

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

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

Answer: D

Explanation:

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

NEW QUESTION 105

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 108

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 +ip4.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 113

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

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

Answer: C

Explanation:

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

NEW QUESTION 118

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

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

Answer: A

Explanation:

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

NEW QUESTION 122

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 125

Which of the following specifications is used to perform disk encryption in a Linux system?

- A. LUKS
- B. TLS
- C. SSL
- D. NFS

Answer: A

Explanation:

LUKS stands for Linux Unified Key Setup, which is a specification for disk encryption on Linux systems. LUKS allows users to encrypt partitions or entire disks using a passphrase or a key file. LUKS also supports multiple keys and key slots, which can be used to unlock the encrypted data. LUKS is compatible with various tools and utilities, such as `cryptsetup`, `dm-crypt`, and `LVM`. References: [How to Encrypt Partitions with LUKS on Linux]

NEW QUESTION 126

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 131

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 136

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

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

Answer: D

Explanation:

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

NEW QUESTION 140

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 142

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 144

A Linux administrator is creating a new sudo profile for the accounting user. Which of the following should be added by the administrator to the sudo configuration file so that the accounting user can run /opt/acc/report as root?

- A. accounting localhost=/opt/acc/report
- B. accounting ALL=/opt/acc/report
- C. %accounting ALL=(ALL) NOPASSWD: /opt/acc/report
- D. accounting /opt/acc/report= (ALL) NOPASSWD: ALL

Answer: C

Explanation:

This answer allows the accounting user to run the /opt/acc/report command as root on any host without entering a password. The % sign indicates that accounting is a group name, not a user name. The ALL keyword means any host, any user, and any command, depending on the context. The NOPASSWD tag overrides the default behavior of sudo, which is to ask for the user's password.

The other answers are incorrect for the following reasons:

? A. accounting localhost=/opt/acc/report

? B. accounting ALL=/opt/acc/report
? D. accounting /opt/acc/report= (ALL) NOPASSWD: ALL

NEW QUESTION 146

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

Output 1:

* httpd.service = The Apache HTTPD Server

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

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

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

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

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

Answer: CD

Explanation:

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

NEW QUESTION 149

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

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

The systems administrator makes additional checks:

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

firewalld is not running
```

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

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

Answer: D

Explanation:

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

The other options are not correct reasons for the firewall rules not being active. iptables is not conflicting with firewalld, because firewalld uses iptables as its backend by default. The wrong system target is not activated, because firewalld is independent of the system target and can be enabled for any target. FIREWALL_ARGS has no value assigned, but this is not a problem, because FIREWALL_ARGS is an optional environment variable that can be used to pass additional arguments to the firewalld daemon, such as `--debug` or `--nofork`. If FIREWALL_ARGS is empty or not defined, firewalld will use its default arguments. References: `firewalld.service(8)` - Linux manual page; `firewall-cmd(1)` - Linux manual page; `systemctl(1)` - Linux manual page

NEW QUESTION 152

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

```
ls -l startup file
```

The following output is returned

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

Which of the following is MOST likely the issue?

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

Answer: A

Explanation:

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

NEW QUESTION 156

An application developer received a file with the following content:

```
##This is a sample Image ## FROM ubuntu:18.04
```

```
MAINTAINER demohut@gmail.com.hac COPY ./app
```

```
RUN make /app
```

```
CMD python /app/app.py RUN apt-get update
```

```
RUN apt-get install -y nginx CMD ["echo","Image created"]
```

The developer must use this information to create a test bed environment and identify the image (myimage) as the first version for testing a new application before moving it to production. Which of the following commands will accomplish this task?

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

Answer: A

Explanation:

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

NEW QUESTION 161

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

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

Answer: B

Explanation:

The command that can be used to check if a partition table is on a disk is `fdisk`. The `fdisk` command can display, create, delete, and modify partitions on a disk. To show the partition table of a disk, the administrator can use `fdisk -l /dev/sdd` (B). References:

? [CompTIA Linux+ Study Guide], Chapter 5: Managing Filesystems and Logical Volumes, Section: Partitioning Disks

? [How to Use Fdisk Command in Linux]

NEW QUESTION 164

An administrator thinks that a package was installed using a snap. Which of the following commands can the administrator use to verify this information?

- A. `snap list`
- B. `snap find`
- C. `snap install`
- D. `snap try`

Answer: A

Explanation:

The `snap list` command is used to display the installed snaps on the system¹. Snaps are self-contained software packages that can be installed and updated across different Linux distributions². The `snap list` command shows the name, version, revision, developer and notes of each snap¹.

The `snap find` command is used to search for snaps in the Snap Store, which is an online repository of snaps². The `snap install` command is used to install snaps from the Snap Store or from a local file². The `snap try` command is used to test a snap without installing it, by mounting a directory that contains the snap files².

These commands are not useful for verifying if a package was installed using a snap.

NEW QUESTION 166

A cloud engineer needs to check the link status of a network interface named `eth1` in a Linux server. Which of the following commands can help to achieve the goal?

- A. ifconfig hw eth1
- B. netstat -r eth1
- C. ss -ti eth1
- D. ip link show eth1

Answer: D

Explanation:

The ip link show eth1 command can be used to check the link status of a network interface named eth1 in a Linux server. It will display information such as the MAC address, MTU, state, and flags of the interface. The ifconfig hw eth1 command is invalid, as hw is not a valid option for ifconfig. The netstat -r eth1 command would display the routing table for eth1, not the link status. The ss -ti eth1 command would display TCP information for sockets associated with eth1, not the link status. References: CompTIA Linux+ (XK0- 005) Certification Study Guide, Chapter 13: Networking Fundamentals, page 436.

NEW QUESTION 168

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

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

Answer: A

Explanation:

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

NEW QUESTION 170

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

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

Answer: A

Explanation:

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

NEW QUESTION 174

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 179

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

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

Answer: D

Explanation:

The environment variable that must be set in remote shell in order to launch the graphical interface is \$DISPLAY. This variable tells X11 applications where to display their windows on screen. It usually has the form hostname:displaynumber.screennumber, where hostname is the name of the computer running the X server, displaynumber is a unique identifier for an X display on that computer, and screennumber is an optional identifier for a screen within an X display. For example, localhost:0.0 means display number 0 on the local host. If the hostname is omitted, it defaults to the local host.

The other options are not correct environment variables for launching the graphical interface. \$RHOST is a variable that stores the name of the remote host, but it is not used by X11 applications. SETENV is a command that sets environment variables in some shells, but it is not an environment variable itself. \$SHELL is a variable that stores the name of the current shell, but it is not related to X11 forwarding. References: How to enable or disable X11 forwarding in an SSH server; How to Configure X11 Forwarding Using SSH In Linux

NEW QUESTION 183

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

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

Answer: A

Explanation:

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

NEW QUESTION 188

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

- A.

```
pvccreate -L1G /dev/app
mkfs.xfs /dev/app
mount /dev/app /opt/app
```
- B.

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

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

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

Answer: D

Explanation:

The list of commands in option D is the correct way to achieve the goal. The commands are as follows:

? fallocate -l 1G /ops/app.img creates a 1GB file named app.img under the /ops directory.

? mkfs.xfs /ops/app.img formats the file as an XFS filesystem.

? mount -o loop /ops/app.img /ops/app mounts the file as a loop device under the /ops/app directory. The other options are incorrect because they either use the wrong commands (dd or truncate instead of fallocate), the wrong options (-t or -f instead of -o), or the wrong order of arguments (/ops/app.img /ops/app instead of /ops/app /ops/app.img). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 323-324.

NEW QUESTION 191

A systems administrator made some changes in the ~/.bashrc file and added an alias command. When the administrator tried to use the alias command, it did not work. Which of the following should be executed FIRST?

- A. source ~/.bashrc
- B. read ~/.bashrc
- C. touch ~/.bashrc
- D. echo ~/.bashrc

Answer: A

Explanation:

The command source ~/.bashrc should be executed first to use the alias command. The source command reads and executes commands from a file in the current shell environment. The ~/.bashrc file is a configuration file that contains commands and aliases that are executed when a new bash shell is started. The administrator made some changes in the ~/.bashrc file and added an alias command, but the changes are not effective until the file is sourced or a new shell is started. The command source ~/.bashrc will reload the file and make the alias command available. The other options are incorrect because they either do not execute the commands in the file (read, touch, or echo) or do not affect the current shell environment (read or echo). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 9: Working with the Linux Shell, page 295.

NEW QUESTION 196

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 198**SIMULATION**

Junior system administrator had trouble installing and running an Apache web server on a Linux server. You have been tasked with installing the Apache web server on the Linux server and resolving the issue that prevented the junior administrator from running Apache.

INSTRUCTIONS

Install Apache and start the service. Verify that the Apache service is running with the defaults.

Typing "help" in the terminal will show a list of relevant event commands.

If at any time you would like to bring back the initial state of the simulation, please click the Reset All button.

CentOS Command Prompt



```
[root@centos7] #
```

- A. Mastered
- B. Not Mastered

Answer: A**Explanation:**

```
yum install httpd
systemctl --now enable httpd systemctl status httpd netstat -tunlp | grep 80
pkill <processname> systemctl restart httpd systemctl status httpd
```

NEW QUESTION 200

Which of the following is a function of a bootloader?

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

Answer: C**Explanation:**

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

NEW QUESTION 202

A systems administrator is notified that the mysqld process stopped unexpectedly. The systems administrator issues the following command: `sudo grep -i -r 'out of memory' /var/log`

The output of the command shows the following:

kernel: Out of memory: Kill process 9112 (mysqld) score 511 or sacrifice child.

Which of the following commands should the systems administrator execute NEXT to troubleshoot this issue? (Select two).

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

Answer: AF**Explanation:**

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

NEW QUESTION 207

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

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

Answer: C

Explanation:

The file that must be updated to ensure the filesystem mounts at boot time is /etc/fstab. This file contains information about the filesystems that are mounted automatically by the mount -a command, which is usually invoked during the system startup. The /etc/fstab file has six fields for each filesystem: device name, mount point, filesystem type, mount options, dump frequency, and pass number. To add a new filesystem to the /etc/fstab file, you need to specify these fields correctly and make sure the mount point directory exists.

The other options are not correct files for controlling persistent mount points of filesystems. The /etc/sysctl file is used to configure kernel parameters at runtime. The /etc/filesystems file is used to specify the order of filesystem types used by mount when no filesystem type is given. The /etc/nfsmount.conf file is used to set options for mounting NFS filesystems. References: Persistently mounting file systems; fstab(5) - Linux manual page

NEW QUESTION 208

A systems administrator created a web server for the company and is required to add a tag for the API so end users can connect. Which of the following would the administrator do to complete this requirement?

- A. hostnamectl status --no-ask-password
- B. hostnamectl set-hostname "\$(perl -le "print" "A" x 86)"
- C. hostnamectl set-hostname Comptia-WebNode -H root@192.168.2.14
- D. hostnamectl set-hostname Comptia-WebNode --transient

Answer: C

Explanation:

The command hostnamectl set-hostname Comptia-WebNode -H root@192.168.2.14 sets the hostname of the web server to Comptia-WebNode and connects to the server using the SSH protocol and the root user. This is the correct way to complete the requirement. The other options are incorrect because they either display the current hostname status (hostnamectl status), set an invalid hostname (hostnamectl set-hostname "\$(perl -le "print" "A" x 86)"), or set a transient hostname that is not persistent (hostnamectl set-hostname Comptia-WebNode --transient). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 9: Managing System Components, page 291.

NEW QUESTION 213

Which of the following commands will display the operating system?

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

Answer: C

Explanation:

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

NEW QUESTION 214

A Linux systems administrator needs to persistently enable IPv4 forwarding in one of the Linux systems. Which of the following commands can be used together to accomplish this task? (Choose two.)

- A. sysctl net.ipv4.ip_forward
- B. sysctl -w net.ipv4.ip_forward=1
- C. echo "net.ipv4.ip_forward=1" >> /etc/sysctl.conf
- D. echo 1 > /proc/sys/net/ipv4/ip_forward
- E. sysctl -p
- F. echo "net.ipv6.conf.all.forwarding=1" >> /etc/sysctl.conf

Answer: BE

Explanation:

The commands that can be used together to persistently enable IPv4 forwarding in one of the Linux systems are sysctl -w net.ipv4.ip_forward=1 and sysctl -p. The first command will use sysctl to write a new value (1) to the net.ipv4.ip_forward kernel parameter, which controls whether IP forwarding is enabled or disabled for IPv4. This will enable IP forwarding immediately without rebooting. However, this change is temporary and will be lost after a reboot or a system reload. To make it permanent, we need to use the second command sysctl -p, which will load kernel parameters from /etc/sysctl.conf file. This file contains key-value pairs of kernel parameters and their values. To make sure that net.ipv4.ip_forward is set to 1 in this file, we can either edit it manually or append it using echo "net.ipv4.ip_forward=1" >> /etc/sysctl.conf.

The other options are not correct commands for persistently enabling IPv4 forwarding. The sysctl net.ipv4.ip_forward command will only display the current value of net.ipv4.ip_forward parameter, but not change it. The echo 1 > /proc/sys/net/ipv4/ip_forward command will write 1 to /proc/sys/net/ipv4/ip_forward file, which is another way to change net.ipv4.ip_forward parameter. However, this change is also temporary and will not survive a reboot or a system reload. The echo "net.ipv6.conf.all.forwarding=1" >> /etc/sysctl.conf command will append a line to /etc/sysctl.conf file that sets net.ipv6.conf.all.forwarding parameter to 1. However,

this parameter controls whether IP forwarding is enabled or disabled for IPv6, not IPv4. References: sysctl(8) - Linux manual page; Configure Linux as a Router (IP Forwarding)

NEW QUESTION 218

A systems administrator has been unable to terminate a process. Which of the following should the administrator use to forcibly stop the process?

- A. kill -1
- B. kill -3
- C. kill -15
- D. kill -HUP
- E. kill -TERM

Answer: E

Explanation:

The administrator should use the command kill -TERM to forcibly stop the process. The kill command is a tool for sending signals to processes on Linux systems. Signals are messages that inform the processes about certain events and actions. The processes can react to the signals by performing predefined or user-defined actions, such as terminating, suspending, resuming, or ignoring. The -TERM option specifies the signal name or number that the kill command should send. The TERM signal, which stands for terminate, is the default signal that the kill command sends if no option is specified. The TERM signal requests the process to terminate gracefully, by closing any open files, releasing any resources, and performing any cleanup tasks. However, if the process does not respond to the TERM signal, the kill command can send a stronger signal, such as the KILL signal, which forces the process to terminate immediately, without any cleanup. The administrator should use the command kill -TERM to forcibly stop the process. This is the correct answer to the question. The other options are incorrect because they either do not terminate the process (kill -1 or kill -3) or do not terminate the process forcibly (kill -15 or kill -HUP). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes, page 431.

NEW QUESTION 221

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

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

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

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

Which of the following commands will permanently resolve the issue?

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

Answer: D

Explanation:

The command that will permanently resolve the issue of high CPU load and average is kill -9 6295. This command will send a SIGKILL signal to the process with the PID 6295, which is the process that is consuming 99.7% of the CPU according to the top output. The SIGKILL signal will terminate the process immediately and free up the CPU resources. The kill command is used to send signals to processes by PID or name.

The other options are not correct commands for resolving this issue. The renice -n -20 6295 command will change the priority (niceness) of the process with PID 6295 to -20, which is the highest priority possible. This will make the process more CPU-intensive, not less. The renice command is used to change the priority of running processes. The pstree -p 6295 command will show a tree of processes with PID 6295 as the root. This will not affect the CPU load or average, but only display information. The pstree command is used to display a tree of processes. The iostat -cy 1 5 command will show CPU and disk I/O statistics for 5 iterations with an interval of 1 second. This will also not affect the CPU load or average, but only display information. The iostat command is used to report CPU and I/O statistics. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Troubleshooting Linux Systems; kill(1) - Linux manual page; renice(1) - Linux manual page; pstree(1) - Linux manual page; iostat(1) - Linux manual page

NEW QUESTION 226

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

XK0-005 Practice Exam Features:

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

100% Actual & Verified — Instant Download, Please Click
[Order The XK0-005 Practice Test Here](#)