

CompTIA

Exam Questions XK0-005

CompTIA Linux+ Certification Exam



NEW QUESTION 1

A systems administrator wants to back up the directory /data and all its contents to /backup/data on a remote server named remote. Which of the following commands will achieve the desired effect?

- A. `scp -p /data remote:/backup/data`
- B. `ssh -i /remote:/backup/ /data`
- C. `rsync -a /data remote:/backup/`
- D. `cp -r /data /remote/backup/`

Answer: C

Explanation:

The command that will back up the directory /data and all its contents to /backup/data on a remote server named remote is `rsync -a /data remote:/backup/`. This command uses the rsync tool, which is a remote and local file synchronization tool. It uses an algorithm to minimize the amount of data copied by only moving the portions of files that have changed. The -a option stands for archive mode, which preserves the permissions, ownership, timestamps, and symbolic links of the files. The /data argument specifies the source directory to be backed up, and the remote:/backup/ argument specifies the destination directory on the remote server. The rsync tool will create a subdirectory named data under /backup/ on the remote server, and copy all the files and subdirectories from /data on the local server.

The other options are not correct commands for backing up a directory to a remote server. The `scp -p /data remote:/backup/data` command will copy the /data directory as a file named data under /backup/ on the remote server, not as a subdirectory with its contents. The -p option preserves the permissions and timestamps of the file, but not the ownership or symbolic links. The `ssh -i /remote:/backup/ /data` command will try to use /remote:/backup/ as an identity file for SSH authentication, which is not valid. The `cp -r /data /remote/backup/` command will try to copy the /data directory to a local directory named /remote/backup/, not to a remote server. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks; rsync(1) - Linux manual page

NEW QUESTION 2

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

```
$ sudo visudo
```

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

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

Which of the following actions will resolve this issue?

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

Answer: B

NEW QUESTION 3

A 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 4

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

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

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

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

Which of the following commands will help unlock the account?

- A. Pam_tally2 --user=dev2 --quiet
- B. pam_tally2 --user=dev2
- C. pam_tally2 --user+dev2 --quiet
- D. pam_tally2 --user=dev2 --reset

Answer: D

Explanation:

To unlock an account that has been locked due to login failures, the administrator can use the command `pam_tally2 --user=dev2 --reset (D)`. This will reset the failure counter for the user “dev2” and allow the user to log in again. The other commands will not unlock the account, but either display or increase the failure count. References:

? [CompTIA Linux+ Study Guide], Chapter 4: Managing Users and Groups, Section: Locking Accounts with pam_tally2

? [How to Lock and Unlock User Account in Linux]

NEW QUESTION 5

After starting an Apache web server, the administrator receives the following error:

```
Apr 23 localhost.localdomain httpd 4618] : (98) Address already in use: AH00072: make_sock: could not bind to address [:]80
```

Which of the following commands should the administrator use to further troubleshoot this issue?

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

Answer: A

Explanation:

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

NEW QUESTION 6

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.org`². This will allow users to resolve the hostnames under `comptia.org` by querying the name server `5.5.5.254`.

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

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

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

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

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

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

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

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

NEW QUESTION 7

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

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

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

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

Answer: AE

Explanation:

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

NEW QUESTION 8

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 9

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

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 10

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

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

# ethtool -S eth0
rx_fw_discards: 4487
```

Which of the following commands will improve the latency issue?

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

Answer: D

Explanation:

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

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

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

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

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

NEW QUESTION 10

A systems engineer has deployed a new application server, but the server cannot communicate with the backend database hostname. The engineer confirms that the application server can ping the database server's IP address. Which of the following is the most likely cause of the issue?

- A. Incorrect DNS servers
- B. Unreachable default gateway
- C. Missing route configuration
- D. Misconfigured subnet mask

Answer: A

Explanation:

This is because the application server can ping the database server's IP address, but not its hostname, which suggests that the DNS resolution is not working properly. DNS servers are responsible for translating hostnames into IP addresses, and vice versa. If the application server has incorrect or unreachable DNS servers configured, it will not be able to resolve the hostname of the database server and communicate with it.

To troubleshoot this issue, the systems engineer should check the DNS configuration on the application server, which is usually stored in the /etc/resolv.conf file.

This file should contain valid nameserver entries that point to the DNS servers that can resolve the database server's hostname. 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.

Alternatively, the systems engineer can use the nslookup or dig commands to test the DNS resolution of the database server's hostname from the application server. These commands will query a specified DNS server and return the IP address of the hostname, or an error message if the resolution fails. For example, to query Google's public DNS server for the IP address of comptia.org, the command would be:

nslookup comptia.org 8.8.8.8 or dig comptia.org @8.8.8.8

NEW QUESTION 11

A systems administrator requires that all files that are created by the user named web have read-only permissions by the owner. Which of the following commands will satisfy this requirement?

- A. chown web:web /home/web
- B. chmod -R 400 /home/web
- C. echo "umask 377" >> /home/web/.bashrc
- D. setfacl read /home/web

Answer: C

Explanation:

The command that will satisfy the requirement of having all files that are created by the user named web have read-only permissions by the owner is echo "umask 377" >> /home/web/.bashrc. This command will append the umask 377 command to the end of the .bashrc file in the web user's home directory. The .bashrc file is a shell script that is executed whenever a new interactive shell session is started by the user. The umask command sets the file mode creation mask, which determines the default permissions for newly created files or directories by subtracting from the maximum permissions (666 for files and 777 for directories). The umask 377 command means that the user does not want to give any permissions to the group or others (3 = 000 in binary), and only wants to give read permission to the owner (7 - 3 = 4 = 100 in binary). Therefore, any new file created by the web user will have read-only permission by the owner (400) and no permission for anyone else. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing Users and Groups; Umask Command in Linux | Linuxize

NEW QUESTION 16

A Linux systems administrator needs to copy files and directories from Server A to Server

- A. Which of the following commands can be used for this purpose? (Select TWO)
- B. rsyslog
 - C. cp
 - D. rsync
 - E. reposync
 - F. scp
 - G. ssh

Answer: CE

Explanation:

The rsync and scp commands can be used to copy files and directories from Server A to Server B. Both commands can use SSH as a secure protocol to transfer data over the network. The rsync command can synchronize files and directories between two locations, using various options to control the copying behavior. The scp command can copy files and directories between two hosts, using similar syntax as cp. The rsyslog command is used to manage system logging, not file copying. The cp command is used to copy files and directories within a single host, not between two hosts. The reposync command is used to synchronize a remote yum repository to a local directory, not copy files and directories between two hosts. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Networking Fundamentals, pages 440-441.

NEW QUESTION 17

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

- A. file filename
- B. touch filename
- C. grep filename
- D. lsof filename

Answer: A

Explanation:

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

NEW QUESTION 20

An engineer needs to insert a character at the end of the current line in the vi text editor. Which of the following will allow the engineer to complete this task?

- A. p
- B. r
- C. bb
- D. A
- E. i

Answer: D

Explanation:

The vi text editor is a popular and powerful tool for editing text files on Linux systems. The vi editor has two modes: command mode and insert mode. In command mode, the user can issue commands to manipulate the text, such as moving the cursor, deleting, copying, pasting, searching, replacing, and saving. In insert mode, the user can type text into the file. To switch from command mode to insert mode, the user can press various keys, such as i, a, o, I, A, or O. To switch from insert mode to command mode, the user can press the Esc key.

To insert a character at the end of the current line in the vi editor, the user can press the A key in command mode. This will move the cursor to the end of the line and switch to insert mode. Then, the user can type the desired character and press Esc to return to command mode. The statement D is correct.

The statements A, B, C, and E are incorrect because they do not perform the desired task. The p key in command mode will paste the previously copied or deleted text after the cursor. The r key in command mode will replace the character under the cursor with another character. The bb key in command mode will move the cursor back two words. The i key in command mode will switch to insert mode before the cursor. References: [How to Use vi Text Editor in Linux]

NEW QUESTION 24

Which of the following can be used as a secure way to access a remote terminal?

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

Answer: B

Explanation:

SSH, or Secure Shell, is a protocol that allows you to access a remote terminal or virtual machine securely over an encrypted connection. You can use SSH to run commands, transfer files, or tunnel network traffic on a remote system. To use SSH, you need an SSH client program on your local system and an SSH server program on the remote system. You also need to authenticate yourself using a username and password or a public/private key pair. SSH is widely used by system administrators, developers, and engineers to remotely manage Linux servers and other devices.

The other options are not correct answers. TFTP, or Trivial File Transfer Protocol, is a simple protocol that allows you to transfer files between systems, but it does not provide any security or encryption features. SCP, or Secure Copy Protocol, is a protocol that uses SSH to securely copy files between systems, but it does not provide a remote terminal access. FTP, or File Transfer Protocol, is another protocol that allows you to transfer files between systems, but it also does not provide any security or encryption features.

NEW QUESTION 26

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

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 29

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

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

The local machine DNS settings are:

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

```
Mail server: mail.example.com
```

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

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

Answer: B

Explanation:

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

NEW QUESTION 30

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

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

Answer: C

Explanation:

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

NEW QUESTION 35

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 37

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

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

Answer: D

Explanation:

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

NEW QUESTION 39

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 44

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 49

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 54

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 58

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 63

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 66

A systems administrator notices the process list on a mission-critical server has a large number of processes that are in state "Z" and marked as "defunct." Which of the following should the administrator do in an attempt to safely remove these entries from the process list?

- A. Kill the process with PID 1.
- B. Kill the PID of the processes.
- C. Kill the parent PID of the processes.
- D. Reboot the server.

Answer: C

Explanation:

As the web search results show, processes in state Z are defunct or zombie processes, which means they have terminated but their parent process has not reaped them properly. They do not consume any resources, but they occupy a slot in the process table. To remove them from the process list, the administrator needs to kill the parent process of the zombies, which will cause them to be reaped by the `init` process (PID 1). Killing the zombies themselves or the `init` process will not have any effect, as they are already dead. Rebooting the server may work, but it is not a safe or efficient option, as it may cause unnecessary downtime or data loss for a mission-critical server.

References

- ? Processes in a Zombie (Z) or Defunct State | Support | SUSE, paragraph 3
- ? linux - Zombie vs Defunct processes? - Stack Overflow, answer by admirableadmin
- ? How To Kill Zombie Processes on Linux | Linux Journal, paragraph 4

NEW QUESTION 71

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 73

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

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

The Linux server has the following system properties CPU: 4 vCPU
Memory: 50GB
Which of the following accurately describes this situation?

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

Answer: A

Explanation:

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

NEW QUESTION 78

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

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

Answer: A

Explanation:

The renice command is used to change the priority of a running process by specifying its PID and the new nice value. The -n flag indicates the amount of change in the nice value, which can be positive or negative. A lower nice value means a higher priority, so -15 will increase the priority of the process with PID 2274. The administrator needs to have root privileges to do this.

References:

? The renice command is listed as one of the commands to manipulate process priority in the web search result 1.

? The renice command is also explained with examples in the web search result 2.

? The CompTIA Linux+ Certification Exam Objectives mention that the candidate should be able to “manage process execution priorities” as part of the System Operation and Maintenance domain1.

NEW QUESTION 79

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 80

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

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

Answer: B

Explanation:

The command docker rm \$(docker ps -aq) will allow the administrator to clean up the containers in an exited state. The docker command is a tool for managing Docker containers on Linux systems. Docker containers are isolated and lightweight environments that can run applications and services without affecting the host system. Docker uses images to create containers, which are files that contain the code, libraries, dependencies, and configuration of the applications and services. The rm option removes one or more containers. The \$(docker ps -aq) is a command substitution that executes the command inside the parentheses and replaces it with the output. The docker ps -aq command lists all the containers, including the ones in an exited state, and shows only their IDs. The docker rm \$(docker ps

-aq) command will remove all the containers, including the ones in an exited state, by passing their IDs to the rm option. This will allow the administrator to clean up the containers in an exited state. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not exist (docker rm --all or docker rm --state exited) or do not remove the containers (docker images prune *). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 571.

NEW QUESTION 82

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

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

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

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

- A. Reboot the serve
- B. The volume will automatically go back to linear mode.
- C. Replace the failed drive and reconfigure the mirror.
- D. Reboot the serve
- E. The volume will revert to stripe mode.
- F. Recreate the logical volume.

Answer: B

Explanation:

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

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

NEW QUESTION 85

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 87

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

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

Answer: D

Explanation:

The lsscsi command can list the SCSI devices on the system, along with their size and device name. The -s option shows the size of each device. The administrator can look for the device that has a size of 500G and note its device name. See lsscsi(8) - Linux man page and How to check Disk Interface Types in

Linux.References1: <https://linux.die.net/man/8/lsscsi>2: <https://www.golinuxcloud.com/check-disk-type-linux/>

NEW QUESTION 88

A systems administrator is trying to track down a rogue process that has a TCP listener on a network interface for remote command-and-control instructions. Which of the following commands should the systems administrator use to generate a list of rogue process names? (Select two).

- A. netstat -antp | grep LISTEN
- B. lsof -iTCP | grep LISTEN
- C. lsof -i:22 | grep TCP
- D. netstat -a | grep TCP
- E. nmap -p1-65535 | grep -i tcp
- F. nmap -sS 0.0.0.0/0

Answer: AB

Explanation:

The best commands to use to generate a list of rogue process names that have a TCP listener on a network interface are A. netstat -antp | grep LISTEN and B. lsof -iTCP | grep LISTEN. These commands will show the process ID (PID) and name of the processes that are listening on TCP ports, which can be used to identify any suspicious or unauthorized processes. The other commands are either not specific enough, not valid, or not relevant for this task. For example:
? C. lsof -i:22 | grep TCP will only show the processes that are listening on port 22, which is typically used for SSH, and not any other ports.
? D. netstat -a | grep TCP will show all the TCP connections, both active and listening, but not the process names or IDs.
? E. nmap -p1-65535 | grep -i tcp will scan all the TCP ports on the local host, but not show the process names or IDs.
? F. nmap -sS 0.0.0.0/0 will perform a stealth scan on the entire internet, which is not only impractical, but also illegal in some countries.

NEW QUESTION 90

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

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

Which of the following commands would resolve the issue?

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 94

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

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

Answer: D

Explanation:

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

NEW QUESTION 97

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

- A. telinit 0
- B. systemctl reboot
- C. systemctl get-default
- D. systemctl emergency

Answer: B

Explanation:

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

NEW QUESTION 101

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 106

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

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 110

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

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

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

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

Answer: BC

Explanation:

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

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

? Create a unit file for the new service in `/etc/systemd/system/` with the name `helpme.service` in the location. This is necessary to register the script as a `systemd` service and enable it to run at startup. A unit file is a configuration file that defines the properties and behavior of a service, such as the description, dependencies, start and stop commands, and environment variables. The unit file should have the extension `.service` and should be placed in the `/etc/systemd/system/` directory. The other option (E) is incorrect because `/etc/init.d` is the directory for `init` scripts, not `systemd` services.

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

NEW QUESTION 112

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 114

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

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

Answer: D

Explanation:

The command `chage -E -1 accountname` will accomplish the task of removing the expiration date of a user account. The `chage` command is a tool for changing user password aging information on Linux systems. The `-E` option sets the expiration date of the user account, and the `-1` value means that the account will never expire. The command `chage -E -1 accountname` will remove the expiration date of the user account named `accountname`. This is the correct command to use to accomplish the task. The

other options are incorrect because they either do not affect the expiration date

(`chgrp`, `passwd`, or `chmod`) or do not exist (`chmod -G`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Users and Groups, page 467.

NEW QUESTION 117

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 119

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

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

Answer: CF

Explanation:

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

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

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

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

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

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

NEW QUESTION 120

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 122

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

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

Answer: A

Explanation:

The Internal Certificate Authority (ICA) is created during the primary Security Management Server installation process. The ICA is a component of Check Point's Public

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

NEW QUESTION 125

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

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

Answer: C

Explanation:

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

NEW QUESTION 126

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 129

Joe, a user, is unable to log in to the Linux system Given the following output:

```
# grep joe /etc/passwd /etc/shadow
/etc/passwd:joe:x:1001:1001::/home/joe:/bin/nologin
/etc/shadow:joe:$6$3uOw6qWx9876jGhgKJsdfH987634534voj.:18883:0:99999:7:::
```

Which of the following command would resolve the issue?

- A. `usermod -s /bin/bash joe`
- B. `pam_tally2 -u joe -r`
- C. `passwd -u joe`
- D. `chage -E 90 joe`

Answer: B

Explanation:

Based on the output of the image sent by the user, Joe is unable to log in to the Linux system because his account has been locked due to too many failed login attempts. The `pam_tally2 -u joe -r` command will resolve this issue by resetting Joe's failed login counter to zero and unlocking his account. This command uses the `pam_tally2` module to manage user account locking based on login failures. The `usermod -s /bin/bash joe` command will change Joe's login shell to `/bin/bash`, but this will not unlock his account. The `passwd -u joe` command will unlock Joe's password if it has been locked by `passwd -l joe`, but this will not reset his failed login counter or unlock his account if it has been locked by `pam_tally2`. The `chage -E 90 joe` command will set Joe's account expiration date to 90 days from today, but this will not unlock his account or reset his failed login counter. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 537.

NEW QUESTION 132

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

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

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 137

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

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

Answer: A

Explanation:

The best command to use to rename the image to match the correct version 2.1.2 is A. `docker tag comptia/app:2.1.1 comptia/app:2.1.2`. This command will create a new tag for the existing image with the new version number, without changing the image content or ID. The other commands are either incorrect or not suitable for this task. For example:

? B. `docker push comptia/app:2.1.1 comptia/app:2.1.2` will try to push two images to a remote repository, but it does not rename the image locally.

? C. `docker rmi comptia/app:2.1.1 comptia/app:2.1.2` will try to remove two images from the local system, but it does not rename the image.

? D. `docker update comptia/app:2.1.1 comptia/app:2.1.2` will try to update the configuration of a running container, but it does not rename the image.

NEW QUESTION 140

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

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

Answer: A

Explanation:

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

? 0: no permission

? 1: execute permission

? 2: write permission

? 4: read permission

? 5: read and execute permissions (4 + 1)

? 6: read and write permissions (4 + 2)

? 7: read, write, and execute permissions (4 + 2 + 1)

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

NEW QUESTION 142

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

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

Answer: A

Explanation:

This file prevents any non-root user from logging in to the system, regardless of the authentication method. The contents of the file are displayed to the user before the login is terminated. This can be useful for system maintenance or security reasons¹².

References: 1: Creating the `/etc/nologin` File - Oracle 2: How to Restrict Log In Capabilities of Users on Ubuntu

NEW QUESTION 145

After installing a new version of a package, a systems administrator notices a new version of the corresponding, service file was installed in order to use the new version of the, service file, which of the following commands must be issued FIRST?

- A. `systemctl status`
- B. `systemctl stop`
- C. `systemctl reinstall`

D. systemctl daemon-reload

Answer: D

Explanation:

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

NEW QUESTION 148

Which of the following commands is used to configure the default permissions for new files?

- A. setenforce
- B. sudo
- C. umask
- D. chmod

Answer: C

Explanation:

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

NEW QUESTION 151

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

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

Answer: B

Explanation:

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

NEW QUESTION 152

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 153

A Linux engineer finds multiple failed login entries in the security log file for application users. The Linux engineer performs a security audit and discovers a security issue. Given the following:

```
# grep -iE '*www*|db' /etc/passwd
www-data:x:502:502:www-data:/var/www:/bin/bash db:x:505:505:db:/opt/db:/bin/bash
```

Which of the following commands would resolve the security issue?

- A. `usermod -d /srv/www-data www-data && usermod -d /var/lib/db db`
- B. `passwd -u www-data && passwd -u db`
- C. `renice -n 1002 -u 502 && renice -n 1005 -u 505`
- D. `chsh -s /bin/false www-data && chsh -s /bin/false db`

Answer: D

Explanation:

This command will use the `chsh` tool to change the login shell of the users `www-data` and `db` to `/bin/false`, which means they will not be able to log in to the system. This will prevent unauthorized access attempts and improve security.

References: 1: Replacing `/bin/bash` with `/bin/false` in `/etc/passwd` file

NEW QUESTION 154

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 156

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

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

Answer: A

Explanation:

The command `useradd -d /comptia/projects user02` will accomplish the task of creating a new user named `user02` with a different home directory.

The `useradd` command is a tool for creating new user accounts on Linux systems. The `-d` option specifies the home directory for the new user, which is the directory where the user's personal files and settings are stored. The `/comptia/projects` is the path of the home directory for the new user, which is different from the default location of `/home/user02`.

The `user02` is the name of the new user. The command `useradd -d /comptia/projects user02` will create a new user named `user02` with a home directory under `/comptia/projects`. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not specify the home directory for the new user (`useradd -m /comptia/projects user02` or `useradd -s /comptia/projects user02`) or do not use the correct option for the home directory (`useradd -b /comptia/projects user02` instead of `useradd -d /comptia/projects user02`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Users and Groups, page 403.

NEW QUESTION 160

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

- A.

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

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


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

D.

```
lvcreate -L 1G -n app app_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 161

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

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

Answer: D

Explanation:

The commands that will list and remove the corresponding packages are dnf history and dnf history undo last. The dnf history command will display a list of all transactions performed by dnf, such as installing, updating, or removing packages. Each transaction has a unique ID, a date and time, an action, and a number of altered packages. The dnf history undo last command will undo the last transaction performed by dnf, meaning that it will reverse all package changes made by that transaction. For example, if the last transaction installed some packages, dnf history undo last will remove them.

The other options are not correct commands for listing and removing corresponding packages. The dnf list command will display a list of available packages in enabled repositories, but not the packages installed by dnf transactions. The dnf remove command will remove specified packages from the system, but not all packages from a specific transaction. The dnf info command will display detailed information about specified packages, but not about dnf transactions. The dnf upgrade command will upgrade all installed packages to their latest versions, but not undo any package changes. References: Handling package management history; dnf(8) - Linux manual page

NEW QUESTION 165

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

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

Answer: C

Explanation:

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

NEW QUESTION 167

A Linux engineer needs to create a custom script, cleanup.sh, to run at boot as part of the system services. Which of the following processes would accomplish this task?

- A. Create a unit file in the /etc/default/ director
- B. systemctl enable cleanupsystemctl is-enabled cleanup
- C. Create a unit file in the /etc/ske1/ director
- D. systemctl enable cleanupsystemctl is-enabled cleanup
- E. Create a unit file in the /etc/systemd/system/ director
- F. systemctl enable cleanupsystemctl is-enabled cleanup
- G. Create a unit file in the /etc/sysctl.d/ director
- H. systemctl enable cleanupsystemctl is-enabled cleanup

Answer: C

Explanation:

The process that will accomplish the task of creating a custom script to run at boot as part of the system services is:

? Create a unit file in the /etc/systemd/system/ directory. A unit file is a configuration

file that defines the properties and behavior of a systemd service. The systemd is a system and service manager that controls the startup and operation of Linux systems. The /etc/systemd/system/ directory is the location where the administrator can create and store custom unit files. The unit file should have a name that matches the name of the script, such as cleanup.service, and should contain the following sections and options:

? Run the command systemctl enable cleanup. This command will enable the service and create the necessary symbolic links to start the service at boot.

? Run the command systemctl is-enabled cleanup. This command will check the status of the service and confirm that it is enabled.

This process will create a custom script, cleanup.sh, to run at boot as part of the system services. This is the correct process to use to accomplish the task. The other options are incorrect because they either use the wrong directory for the unit file (/etc/default/, /etc/skel/, or /etc/sysctl.d/) or do not create a unit file at all. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing System Services, pages 457-459.

NEW QUESTION 168

Which of the following should be used to verify the integrity of a file?

- A. sha256sum
- B. fsck
- C. gpg —d
- D. hashcat

Answer: A

Explanation:

The best tool to use to verify the integrity of a file is A. sha256sum. This tool will compute and display the SHA-256 hash of a file, which is a 64-digit hexadecimal number that uniquely identifies the file's content. By comparing the hash of a downloaded file with the hash provided by the file owner or source, you can confirm that the file has not been altered or corrupted during the transfer. The other tools are either not relevant or not suitable for this task. For example:

? B. fsck is a tool for checking and repairing the file system, but it does not verify the integrity of individual files.

? C. gpg -d is a tool for decrypting files that have been encrypted with GnuPG, but it does not verify the integrity of unencrypted files.

? D. hashcat is a tool for cracking passwords or hashes, but it does not verify the integrity of files.

NEW QUESTION 172

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 175

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 177

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

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

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

Which of the following commands will resolve this issue?

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

Answer: C

Explanation:

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

NEW QUESTION 182

A systems administrator is investigating why one of the servers has stopped connecting to the internet.

```
#curl http://google.com
curl: (6) Could not resolve host: google.com

#cat /etc/resolv.conf
search user.company.com company.com
#nameserver 10.10.10.10

#ip route
0.0.0.0/0 via 10.0.5.1 dev eth0 proto static metric 100
10.0.0.0/16 dev eth0 proto kernel scope link src 10.0.3.60 metric 101

#nmcli connection show
NAME                UUID                                  TYPE      DEVICE
eth0                 ba4a3d30-efdc-4fa5-83d3-3721fd4aff75  ethernet  eth0
Wired connection 1   8d569d5a-22a2-356d-8532-9a2638f11b5a5  ethernet  --
```

Which of the following is causing the issue?

- A. The DNS address has been commented out in the configuration file.
- B. The search entry in the /etc/resolv.conf file is incorrect.
- C. Wired connection 1 is offline.
- D. No default route is defined.

Answer: D

Explanation:

The issue is caused by the lack of a default route defined in the /etc/sysconfig/network-scripts/ifcfg-enp0s3 file. A default route is a special route that specifies where to send packets that do not match any other routes in the routing table. Without a default route, the server will not be able to communicate with hosts outside its local network. The default route is usually configured with the GATEWAY option in the network interface configuration file. For example, to set the

default gateway to 192.168.1.1, the file should contain:
GATEWAY=192.168.1.1

The other options are not causing the issue. The DNS address is not commented out in the configuration file, it is specified with the DNS1 option. The search entry in the /etc/resolv.conf file is correct, it specifies the domain name to append to unqualified hostnames. Wired connection 1 is online, as indicated by the ONBOOT=yes option and the output of ip link show enp0s3 command. References: Configuring IP Networking with nmcli; Configuring IP Networking with ifcfg Files

NEW QUESTION 187

A Linux engineer is setting the sticky bit on a directory called devops with 755 file permission. Which of the following commands will accomplish this task?

- A. chown -s 755 devops
- B. chown 1755 devops
- C. chmod -s 755 devops
- D. chmod 1755 devops

Answer: D

Explanation:

The command that will set the sticky bit on a directory called devops with 755 file permission is chmod 1755 devops. This command will use chmod to change the mode of the directory devops to 1755, which means that the owner has read, write, and execute permissions (7), the group has read and execute permissions (5), and others have read and execute permissions (5). The first digit 1 indicates that the sticky bit is set on the directory, which is a special permission that prevents users from deleting or renaming files in the directory that they do not own.

The other options are not correct commands for setting the sticky bit on a directory. The chown -s 755 devops command is invalid because chown is used to change the owner and group of files or directories, not their permissions. The -s option for chown is used to remove a symbolic link, not to set the sticky bit. The chown 1755 devops command is also invalid because chown does not accept numeric arguments for changing permissions. The chmod -s 755 devops command will remove the sticky bit from the directory devops, not set it. References: chmod(1) - Linux manual page; How to Use SUID, SGID, and Sticky Bits on Linux

NEW QUESTION 189

A systems administrator is investigating a service that is not starting up. Given the following information:

```
root@localhost ~]# systemctl status network
network.service - LSB: Bring up/down networking
Loaded: loaded (/etc/rc.d/init.d/network; bad; vendor preset: disabled)
Active: failed (Result: exit-code) since Jan 2022-01-02 22:55:15 CST;
Docs: man:systemd-sysv-generator(8)
Process: 1083 ExecStart=/etc/rc.d/init.d/network start (code=exited, status=1/FAILURE)
Jan 02 22:55:15 localhost.localdomain network[1083]: Bringing up interface enp0s25: Error: Con...n.
Jan 02 22:55:15 localhost.localdomain network[1083]: [FAILED]
[...]
```

Which of the following systemd commands should the administrator use in order to obtain more details about the failing service?

- A. systemctl analyze network
- B. systemctl info network
- C. sysctl -a network
- D. journalctl -xu network

Answer: D

Explanation:

The systemd is a system and service manager for Linux systems that provides a standard way to control and monitor system services. The systemd uses various commands and tools to manage and troubleshoot system services, such as systemctl, sysctl, and journalctl. The systemctl command is used to start, stop, enable, disable, restart, reload, status, and list system services. The sysctl command is used to configure kernel parameters at runtime. The journalctl command is used to view and filter the logs of system services.

To investigate a service that is not starting up, the administrator can use the journalctl command with the -xu option. The -x option enables verbose output that includes explanatory text and priority information. The -u option filters the output by a specific unit name, such as network.service. Therefore, the command journalctl -xu network will show detailed logs of the network service, which can help identify the cause of the failure. The statement D is correct.

The statements A, B, and C are incorrect because they do not provide more details about the failing service. The systemctl analyze network command does not exist.

The systemctl info network command shows basic information about the network unit, such as description, load state, active state, sub state, and main PID. The sysctl -a network command shows all kernel parameters related to network settings. References: [How to Use Systemd to Manage System Services]

NEW QUESTION 191

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

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

Answer: D

Explanation:

The yum history undo last command will undo the last transaction, which in this case is the installation of the httpd RPM package and its dependencies. This will remove the packages that were installed and restore the previous state of the system. See How to undo or redo yum transactions and yum history. References: 1: <https://www.redhat.com/sysadmin/undo-redo-yum-transactions> 2: <https://man7.org/linux/man-pages/man8/yum.8.html#HISTORY>

NEW QUESTION 196

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions 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](#)