



CompTIA

Exam Questions XK0-005

CompTIA Linux+ Certification Exam

NEW QUESTION 1

A Linux administrator was notified that a virtual server has an I/O bottleneck. The Linux administrator analyzes the following output:

```
root@linux:~# uptime
18:43:47 up 1 day, 19:58, 1 user, load average: 9.90, 5.83, 2.49
root@linux:~# vmstat 10 10
procs -----memory----- --swap----- --io--- -system- -----cpu-----

 r b swpd   free   buff   cache  si    so bi    bo    in    cs us  sy  id  wa  st
 13 0 5520 141228 98932 2325312 0      2 10    28   192   167  1  0 99  0  0
 10 0 5608 131280 98932 2325324 0 26211 0 26211 342   393 91  9  0  0  0
 10 0 5528   1096 98932 2325324 0  5242 0  5242 333   402 96  4  0  0  0

root@linux:~# free -m
              total    used     free shared buff/cache   available
Mem:           3933    1454       110      33        2368        2202
Swap:           1497         5       1491
```

Given there is a single CPU in the sever, which of the following is causing the slowness?

- A. The system is running out of swap space.
- B. The CPU is overloaded.
- C. The memory is exhausted.
- D. The processes are paging.

Answer: B

Explanation:

The slowness is caused by the CPU being overloaded. The iostat command shows that the CPU utilization is 100%, which means that there are more processes competing for CPU time than the CPU can handle. The other options are incorrect because:
 ? The system is not running out of swap space, as shown by the iostat command, which shows that there is no swap activity (si and so columns are zero).
 ? The memory is not exhausted, as shown by the free -m command, which shows that there is still available memory (avail column) and free buffer/cache memory (buff/cache column).
 ? The processes are not paging, as shown by the vmstat command, which shows that there are no major page faults (majflt column) and no swap activity (si and so columns). References: CompTIA Linux+ Study Guide, Fourth Edition, page 417- 419, 424-425.

NEW QUESTION 2

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 3

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

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

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

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

Answer: A

Explanation:

The command vmstat will most likely be run next by the administrator to troubleshoot the system performance. The vmstat command is a tool for reporting virtual memory statistics on Linux systems. The command shows information about processes, memory, paging, block IO, interrupts, and CPU activity. The command can help the administrator identify the source of the performance issue, such as high CPU usage, low free memory, excessive swapping, or disk IO bottlenecks. The

command can also be used with an interval and a count to display the statistics repeatedly over time and observe the changes. The command `vmstat` will provide useful information for diagnosing the system performance and finding the root cause of the issue. This is the most likely command to run next after the `top` command. The other options are incorrect because they either do not show the virtual memory statistics (`strace` or `lsof`) or do not provide more information than the `top` command (`htop`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, page 425.

NEW QUESTION 4

The administrator `comptia` is not able to perform privileged functions on a newly deployed system. Given the following command outputs:

```
[root@newserver ~]# id comptia
uid=1000(comptia) gid=1000(comptia) groups=1000(comptia)

[root@newserver ~]# cat /etc/sudoers.d/admin
%admin ALL= (root) NOPASSWD: EXEC: /usr/bin/ps, /usr/bin/chmod, /usr/bin/yum, /usr/bin/cat, /usr/sbin/lvm,
/usr/sbin/pvs

[root@newserver ~]# grep comptia /etc/passwd
comptia:x:1000:1000:comptia:/home/comptia:/bin/bash

[root@newserver ~]# chage -l comptia
Last password change : never
Password expires : never
Password inactive : never
Account expires : never
Minimum number of days between password change : 0
Maximum number of days between password change : 99999
Number of days of warning before password expires : 7
```

Which of the following is the reason that the administrator is unable to perform the assigned duties?

- A. The administrator needs a password reset.
- B. The administrator is not a part of the correct group.
- C. The administrator did not update the sudo database.
- D. The administrator's credentials need to be more complex.

Answer: B

Explanation:

The reason that the administrator is unable to perform the assigned duties is because the administrator is not a part of the correct group. This is option B. Based on the image that you sent, I can see that the user `comptia` has a user ID and a group ID of 1000, and belongs to only one group, which is also `comptia`. However, the `sudoers` file, which defines the permissions for users to run commands as root or other users, does not include the `comptia` group in any of the entries. Therefore, the user `comptia` cannot use `sudo` to perform privileged functions on the system.

The other options are incorrect because:

* A. The administrator needs a password reset.

This is not true, because the password aging information for the user `comptia` shows that the password was last changed on Oct 24, 2023, and it does not expire until Jan 22, 2024. There is no indication that the password is locked or expired.

* C. The administrator did not update the sudo database.

This is not necessary, because the sudo database is automatically updated whenever the `sudoers` file is modified. There is no separate command to update the sudo database.

* D. The administrator's credentials need to be more complex.

This is not relevant, because the complexity of the credentials does not affect the ability to use `sudo`. The `sudoers` file does not specify any password policy for the users or groups that are allowed to use `sudo`.

NEW QUESTION 5

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

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

Answer: D

Explanation:

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

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

References: `rpm(8)` - Linux manual

page; Using RPM to Verify Installed Packages

NEW QUESTION 6

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

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

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

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

Which of the following commands will help unlock the account?

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 7

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 8

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

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

Answer: D

Explanation:

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

NEW QUESTION 9

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

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

Answer: CF

Explanation:

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

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

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

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

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

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

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

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

NEW QUESTION 10

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

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

Answer: D

Explanation:

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

The other options are incorrect because:

* A. Run `git exclude app.conf`

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

* B. Run `git stash app.conf`

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

* C. Add `app.conf` to `.exclude`

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

References:

? Git - `gitignore` Documentation

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

? Ignoring files - GitHub Docs

? [CompTIA Linux+ Certification Exam Objectives]

NEW QUESTION 10

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

```
[root@system] # cat mydocs.mount [Unit]
```

```
Description=Mount point for My Documents drive [Mount]
```

```
What=/dev/drv/disk/by-uuid/94afc9b2-ac34-ccff-88ae-297ab3c7ff34 Where=/home/user1/My Documents
```

```
Options=defaults Type=xfs
```

```
[Install]
```

```
WantedBy=multi-user.target
```

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

- A. Rename the mount file to `home-user1-My\Documents.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\Documents`.
- 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. ReferencesThe 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 11

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 14

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

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

Answer: D

Explanation:

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

? gedit & disown will launch a graphical text editor in the background and detach it from the terminal, but it will not execute any service.

? kill 9 %1 will terminate the job with ID 1 using a SIGKILL signal, which cannot be ignored or handled by the process.

? fg %1 will bring the job with ID 1 to the foreground, where it will occupy the terminal until it finishes or is stopped. References: CompTIA Linux+ Study Guide, Fourth Edition, page 181-182.

NEW QUESTION 16

In order to copy data from another VLAN, a systems administrator wants to temporarily assign IP address 10.0.6.5/24 to the newly added network interface enp1s0f1. Which of the following commands should the administrator run to achieve the goal?

- A. ip addr add 10.0.6.5/24 dev enp1s0f1
- B. echo "IPV4_ADDRESS=10.0.6.5/24" > /etc/sysconfig/network-scripts/ifcfg-enp1s0f1
- C. ifconfig 10.0.6.5/24 enp1s0f1
- D. nmcli conn add ipv4.address-10.0.6.5/24 ifname enp1s0f1

Answer: A

Explanation:

The command ip addr add 10.0.6.5/24 dev enp1s0f1 will achieve the goal of temporarily assigning IP address 10.0.6.5/24 to the newly added network interface enp1s0f1. The ip command is a tool for managing network interfaces and routing on Linux systems. The addr option specifies the address manipulation mode. The add option adds a new address to an interface. The 10.0.6.5/24 is the IP address and the subnet mask in CIDR notation. The dev option specifies the device name. The enp1s0f1 is the name of the network interface. The command ip addr add 10.0.6.5/24 dev enp1s0f1 will add the IP address 10.0.6.5/24 to the network interface enp1s0f1, which will allow the administrator to copy data from another VLAN. This is the correct command to use to achieve the goal. The other options are incorrect because they either do not add a new address to an interface (echo "IPV4_ADDRESS=10.0.6.5/24" > /etc/sysconfig/network-scripts/ifcfg-enp1s0f1 or ifconfig 10.0.6.5/24 enp1s0f1) or do not use the correct syntax for the command (nmcli conn add ipv4.address-10.0.6.5/24 ifname enp1s0f1 instead of nmcli conn add type ethernet ipv4.address 10.0.6.5/24 ifname enp1s0f1). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 385.

NEW QUESTION 19

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

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 22

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

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

Answer: A

Explanation:

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

NEW QUESTION 26

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

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

Answer: B

Explanation:

The administrator should set PermitRootLogin to no in the /etc/ssh/sshd_config file to remove the possibility of remote administrative login via the SSH service. The PermitRootLogin directive controls whether the root user can log in using SSH. Setting it to no will deny any remote login attempts by the root user. This will

harden the server and prevent unauthorized access. The administrator should also restart the sshd service after making the change. The other options are incorrect because they either do not affect the SSH service (/etc/hosts.deny or /etc/pam.d/sshd) or do not prevent remote administrative login (PubKeyAuthentication). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 413.

NEW QUESTION 27

User1 is a member of the accounting group. Members of this group need to be able to execute but not make changes to a script maintained by User2. The script should not be accessible to other users or groups. Which of the following will give proper access to the script?

- A. `chown user2:accounting script.sh chmod 750 script.sh`
- B. `chown user1:accounting script.sh chmod 777 script.sh`
- C. `chown accounting:user1 script.sh chmod 057 script.sh`
- D. `chown user2:accounting script.sh chmod u+x script.sh`

Answer: A

Explanation:

The commands that will give proper access to the script are:

? `chown user2:accounting script.sh`: This command will change the ownership of the script to user2 as the owner and accounting as the group. The `chown` command is a tool for changing the owner and group of files and directories on Linux systems. The `user2:accounting` is the user and group name that the command should assign to the script. The `script.sh` is the name of the script that the command should modify. The command `chown user2:accounting script.sh` will ensure that user2 is the owner of the script and accounting is the group of the script, which will allow user2 to maintain the script and the accounting group to access the script.

? `chmod 750 script.sh`: This command will change the permissions of the script to 750, which means read, write, and execute for the owner; read and execute for the group; and no access for others. The `chmod` command is a tool for changing the permissions of files and directories on Linux systems. The permissions are represented by three digits in octal notation, where each digit corresponds to the owner, group, and others. Each digit can have a value from 0 to 7, where each value represents a combination of read, write, and execute permissions. The 750 is the permission value that the command should assign to the script.

The `script.sh` is the name of the script that the command should modify. The command `chmod 750 script.sh` will ensure that only the owner and the group can execute the script, but not make changes to it, and that the script is not accessible to other users or groups.

The commands that will give proper access to the script are `chown user2:accounting script.sh` and `chmod 750 script.sh`. This is the correct answer to the question.

The other options are incorrect because they either do not give proper access to the script (`chown user1:accounting script.sh` or `chown accounting:user1 script.sh`) or do not change the permissions of the script (`chmod 777 script.sh` or `chmod u+x script.sh`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing File Permissions and Ownership, pages 346-348.

NEW QUESTION 31

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

`ping: comptia.org: Name or service not known`

Which of the following files should the administrator verify?

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

Answer: C

Explanation:

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

```
nameserver 8.8.8.8 nameserver 8.8.4.4
```

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

NEW QUESTION 36

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

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

Answer: D

Explanation:

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

NEW QUESTION 38

Which of the following directories is the mount point in a UEFI system?

- A. `/sys/efi`
- B. `/boot/efi`
- C. `/efi`

D. /etc/efi

Answer: B

Explanation:

The /boot/efi directory is the mount point in a UEFI system. This directory contains the EFI System Partition (ESP), which stores boot loaders and other files required by UEFI firmware. The /sys/efi directory does not exist by default in Linux systems. The /efi directory does not exist by default in Linux systems. The /etc/efi directory does not exist by default in Linux systems. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing the Linux Boot Process, page 398.

NEW QUESTION 42

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

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

Answer: C

Explanation:

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

NEW QUESTION 43

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

- A. visudo -c
- B. test -f /etc/sudoers
- C. sudo vi check
- D. cat /etc/sudoers | tee test

Answer: A

Explanation:

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

NEW QUESTION 48

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 49

While inspecting a recently compromised Linux system, the administrator identified a number of processes that should not have been running:

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
5545	joe	30	-10	5465	56465	8254	R	0.5	1.5	00:35.3	upload.sh
2567	joe	30	-10	6433	75544	9453	R	0.7	1.8	00:25.1	upload_passwd.sh
8634	joe	30	-10	3584	74537	6435	R	0.3	1.1	00:17.6	uploadpw.sh
4846	joe	30	-10	6426	63234	9683	R	0.8	1.9	00:22.2	upload_shadow.sh

Which of the following commands should the administrator use to terminate all of the identified processes?

- A. pkill -9 -f "upload*.sh"

- B. kill -9 "upload*.sh"
- C. killall -9 -upload*.sh"
- D. skill -9 "upload*.sh"

Answer: A

Explanation:

The pkill -9 -f "upload*.sh" command will terminate all of the identified processes. This command will send a SIGKILL signal (-9) to all processes whose full command line matches the pattern "upload*.sh" (-f). This signal will force the processes to terminate immediately without giving them a chance to clean up or save their state. The kill -9 "upload*.sh" command is invalid, as kill requires a process ID (PID), not a pattern. The killall -9 "upload*.sh" command is incorrect, as killall requires an exact process name, not a pattern. The skill -9 "upload*.sh" command is incorrect, as skill requires a username or a session ID (SID), not a pattern. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Memory and Process Execution, page 470.

NEW QUESTION 54

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

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

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

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

Answer: C

Explanation:

The command that can be used to resolve the issue of being unable to remove a particular data folder is chattr -R -i data/. This command will use the chattr utility to change file attributes on a Linux file system. The -R option means that chattr will recursively change attributes of directories and their contents. The -i option means that chattr will remove (unset) the immutable attribute from files or directories. When a file or directory has the immutable attribute set, it cannot be modified, deleted, or renamed.

The other options are not correct commands for resolving this issue. The chgrp -R 755 data/ command will change the group ownership of data/ and its contents recursively to 755, which is not a valid group name. The chgrp command is used to change group ownership of files or directories. The chmod -R 777 data/ command will change the file mode bits of data/ and its contents recursively to 777, which means that everyone can read, write, and execute them. However, this will not remove the immutable attribute, which prevents deletion or modification regardless of permissions. The chmod command is used to change file mode bits of files or directories. The chown -R data/ command is incomplete and will produce an error. The chown command is used to change the user and/or group ownership of files or directories, but it requires at least one argument besides the file name. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 7: Managing Disk Storage; chattr(1) - Linux manual page; chgrp(1) - Linux manual page; chmod(1) - Linux manual page; chown(1) - Linux manual page

NEW QUESTION 59

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

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

Answer: D

Explanation:

To delete a file from a Git repository, the administrator can use the command git rm app.conf (D). This will remove the file "app.conf" from the working directory and stage it for deletion from the repository. The administrator can then commit the change with git commit -m "Delete app.conf" to finalize the deletion. The other commands will not delete the file, but either tag, commit, or checkout the file. References:

? [CompTIA Linux+ Study Guide], Chapter 10: Working with Git, Section: Deleting Files with Git
? [How to Delete Files from Git]

NEW QUESTION 63

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

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

Answer: C

Explanation:

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

NEW QUESTION 67

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 69

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

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

Given the following analysis:

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

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

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

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

Which of the following is causing the error message?

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

Answer: D

Explanation:

File attributes are preventing file modification, which is causing the error message. The output of `lsattr /home/user/data.txt` shows that the file has the immutable

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

NEW QUESTION 73

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

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

Answer: B

Explanation:

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

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

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

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

References:

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

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

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

NEW QUESTION 76

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 77

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 79

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 84

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

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

Answer: C

Explanation:

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

NEW QUESTION 86

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

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

Answer: B

Explanation:

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

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

? [How to Use .gitignore File]

NEW QUESTION 91

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 93

A systems administrator is tasked with installing GRUB on the legacy MBR of the SATA hard drive. Which of the following commands will help the administrator accomplish this task?

- A. grub-install /dev/hda
- B. grub-install /dev/sda
- C. grub-install /dev/sr0
- D. grub-install /dev/hd0,0

Answer: B

Explanation:

The command that will help the administrator install GRUB on the legacy MBR of the SATA hard drive is grub-install /dev/sda. This command will install GRUB on the master boot record (MBR) of the first SATA disk (/dev/sda). The MBR is the first sector of a disk that contains boot code and a partition table. GRUB will overwrite the boot code and place its own code that can load GRUB modules and configuration files from a specific partition. The other options are not correct commands for installing GRUB on the legacy MBR of the SATA hard drive. The grub-install /dev/hda command will try to install GRUB on the first IDE disk (/dev/hda), which may not exist or may not be bootable. The grub-install /dev/sr0 command will try to install GRUB on the first SCSI CD-ROM device (/dev/sr0), which is not a hard drive and may not be bootable. The grub-install /dev/hd0,0 command is invalid because grub-install does not accept partition names as arguments, only disk names. References: Installing GRUB using grub-install; GRUB Manual

NEW QUESTION 95

A Linux administrator is adding a new configuration file to a Git repository. Which of the following describes the correct order of Git commands to accomplish the task successfully?

- A. pull -> push -> add -> checkout
- B. pull -> add -> commit -> push
- C. checkout -> push -> add -> pull
- D. pull -> add -> push -> commit

Answer: B

Explanation:

The correct order of Git commands to add a new configuration file to a Git repository is pull -> add -> commit -> push. The pull command will fetch and merge the changes from the remote repository to the local repository, ensuring that the local repository is up to date. The add command will stage the new configuration file for the next commit, marking it as a new file to be tracked by Git. The commit command will create a new snapshot of the project state with the new configuration file and a descriptive message. The push command will publish the commit to the remote repository, updating the remote branch with the new configuration file. The pull -> push -> add -> checkout order is incorrect, as it will not create a commit for the new configuration file, and it will switch to a different branch without pushing the changes. The checkout -> push -> add -> pull order is incorrect, as it will switch to a different branch before adding the new configuration file, and it will overwrite the local changes with the remote changes without creating a commit. The pull -> add -> push -> commit order is incorrect, as it will not create a commit before pushing the changes, and it will create a commit that is not synchronized with the remote branch. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 20: Writing and Executing Bash Shell Scripts, page 579.

NEW QUESTION 98

A systems administrator is adding a Linux-based server and removing a Windows-based server from a cloud-based environment. The changes need to be validated before they are applied to the cloud-based environment. Which of the following tools should be used to meet this requirement?

- A. Ansible
- B. git clone
- C. git pull
- D. terraform plan

Answer: D

Explanation:

Terraform is a tool for building, changing, and managing infrastructure as code in a cloud-based environment. Terraform uses configuration files to describe the desired state of the infrastructure and applies changes accordingly. Terraform supports various cloud providers, such as AWS, Azure, Google Cloud Platform, and more.

To validate changes before they are applied to the cloud-based environment, the administrator can use the terraform plan command. This command will compare the current state of the infrastructure with the desired state defined in the configuration files and show what actions will be performed to achieve the desired state. This command will not make any changes to the infrastructure but only show a plan of changes. The statement D is correct.

The statements A, B, and C are incorrect because they do not validate changes before they are applied to the cloud-based environment. Ansible is another tool for automating infrastructure management, but it does not have a plan command. Git clone and git pull are commands for working with git repositories, which are used for version control of code. References: [How to Use Terraform to Manage Cloud Infrastructure]

NEW QUESTION 103

A Linux systems administrator is setting up a new web server and getting 404 - NOT FOUND errors while trying to access the web server pages from the browser. While working on the diagnosis of this issue, the Linux systems administrator executes the following commands:

```
# getenforce
Enforcing

# matchpathcon -V /var/www/html/*
/var/www/html/index.html has context unconfined_u:object_r:user_home_t:s0, should be system_u:object_r:httpd_sys_content_t:s0
/var/www/html/page1.html has context unconfined_u:object_r:user_home_t:s0, should be system_u:object_r:httpd_sys_content_t:s0
```

Which of the following commands will BEST resolve this issue?

- A. sed -i 's/SELINUX=enforcing/SELINUX=disabled/' /etc/selinux/config
- B. restorecon -R -v /var/www/html
- C. setenforce 0
- D. setsebool -P httpd_can_network_connect_db on

Answer: B

Explanation:

The command restorecon -R -v /var/www/html will best resolve the issue. The issue is caused by the incorrect SELinux context of the web server files under the /var/www/html directory. The output of ls -Z /var/www/html shows that the files have the type user_home_t, which is not allowed for web content. The command restorecon restores the default SELinux context of files based on the policy rules. The options -R and -v are used to apply the command recursively and verbosely. This command will change the type

of the files to httpd_sys_content_t, which is the correct type for web content. This will allow the web server to access the files and serve the pages to the browser. The other options are incorrect because they either disable SELinux entirely (sed -i 's/SELINUX=enforcing/SELINUX=disabled/' /etc/selinux/config or setenforce 0), which is not a good security practice, or enable an unnecessary boolean (setsebool -P httpd_can_network_connect_db on), which is not related to the issue.

References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 535.

NEW QUESTION 107

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

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

Answer: A

Explanation:

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

NEW QUESTION 110

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

```
$ nmap 192.168.10.36
```

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

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

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

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

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

Answer: A

Explanation:

This is because the port 22/tcp is shown as filtered by nmap, which means that nmap cannot determine whether the port is open or closed because a firewall or other device is blocking its probes. If the SSH server was not running on the remote server, the port would be shown as closed, which means that nmap received a TCP RST packet in response to its probe. If the remote SSH server was using SSH protocol version 1, the port would be shown as open, which means that nmap received a TCP SYN/ACK packet in response to its probe. If the SSH host key on the remote server had expired, the port would also be

shown as open, but the SSH client would display a warning message about the host key verification failure. Therefore, the best explanation for the filtered state of the port 22/tcp is that a firewall is preventing nmap from reaching the SSH server.

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

? Nmap scan what does STATE=filtered mean?

? How to find ports marked as filtered by nmap

? Technical Tip: NMAP scan shows ports as filtered

NEW QUESTION 112

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 115

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

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

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

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

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

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

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

Answer: C

Explanation:

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

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

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

The other options are incorrect because:

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

This is not true, because the output for the first command shows that the /ftpusers/ filesystem has 26% of disk space available, which means that there is enough space for users to create files. The permissions of the files and directories are not shown in the output, but they are not relevant to the issue of inode exhaustion.

* B. The ftpusers filesystem does not have enough space.

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

* D. ftpusers is mounted as read only.

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

NEW QUESTION 116

An administrator deployed a Linux server that is running a web application on port 6379/tcp.

SELinux is in enforcing mode based on organization policies. The port is open on the firewall.

Users who are trying to connect to a local instance of the web application receive Error 13, Permission denied. The administrator ran some commands that resulted in the following output:

```
# semanage port -l | egrep '(^http_port_t|6379) '
http_port_t tcp 80, 81, 443, 488, 8008, 8009, 8443, 9000

# curl http://localhost/App.php
Cannot connect to App Server.
```

Which of the following commands should be used to resolve the issue?

- A. `semanage port -d -t http_port_t -p tcp 6379`
- B. `semanage port -a -t http_port_t -p tcp 6379`
- C. `semanage port -a http_port_t -p top 6379`
- D. `semanage port -l -t http_port_tcp 6379`

Answer: B

Explanation:

The command `semanage port -a -t http_port_t -p tcp 6379` adds a new port definition to the SELinux policy and assigns the type `http_port_t` to the port `6379/tcp`. This allows the web application to run on this port and accept connections from users. This is the correct way to resolve the issue. The other options are incorrect because they either delete a port definition (`-d`), use the wrong protocol (`top` instead of `tcp`), or list the existing port definitions (`-l`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 535.

NEW QUESTION 120

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 121

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 123

A systems administrator was tasked with assigning the temporary IP address/netmask `192.168.168.1/255.255.255.255` to the interface `eth0` of a Linux server.

When adding the address, the following error appears:

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

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

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

- A. The CIDR value `/33` should be `/32` instead.

- B. There is no route to 192.168.168.1/33.
- C. The interface eth0 does not exist.
- D. The IP address 192.168.168.1 is already in use.

Answer: A

Explanation:

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

NEW QUESTION 125

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

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

Answer: B

Explanation:

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

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

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

NEW QUESTION 128

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 130

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

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

Answer: B

Explanation:

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

NEW QUESTION 131

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 132

Which of the following files holds the system configuration for journal when running systemd?

- A. /etc/systemd/journald.conf
- B. /etc/systemd/systemd-journalctl.conf
- C. /usr/lib/systemd/journalctl.conf
- D. /etc/systemd/systemd-journald.conf

Answer: A

Explanation:

The file that holds the system configuration for journal when running systemd is /etc/systemd/journald.conf. This file contains various settings that control the behavior of the journald daemon, which is responsible for collecting and storing log messages from various sources. The journald.conf file can be edited to change the default values of these settings, such as the storage location, size limits, compression, and forwarding options of the journal files. The file also supports a drop-in directory /etc/systemd/journald.conf.d/ where additional configuration files can be placed to override or extend the main file. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks; journald.conf(5) - Linux manual page

NEW QUESTION 134

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 135

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

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

Answer: A

Explanation:

The /sbin/nologin shell is a special shell that prevents the user from logging into an interactive session. It is commonly used for system accounts that are not meant to be accessed by users, such as daemon or service accounts. When a user tries to log in with this shell, they will see a message like "This account is currently not available" and the login will fail.

References:

- ? The /sbin/nologin shell is listed as one of the valid shells in the /etc/shells file1.
- ? The CompTIA Linux+ Certification Exam Objectives mention that the candidate should be able to "configure and manage system accounts and groups, including password aging and restricted shells" as part of the Hardware and System Configuration domain2.
- ? The usermod command can be used to change the user's login shell with the -s or --shell option3. For example, to change the shell of a user named daemon to /sbin/nologin, the command would be: sudo usermod -s /sbin/nologin daemon

NEW QUESTION 138

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 142

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 143

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

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

Answer: A

Explanation:

A long-running process is a program that takes a long time to complete or runs indefinitely on a Linux system. A foreground process is a process that runs in the current terminal and receives input from the keyboard and output to the screen. A background process is a process that runs in the background and does not interact with the terminal. A background process can continue running even if the terminal is closed or disconnected.

To start a long-running process in the background, the user can append an ampersand (`&`)

to the command, such as `someapp &`. This will run `someapp` in the background and return control to the terminal immediately.

To move a long-running process from the foreground to the background, the user can use two keystrokes: `Ctrl+Z` and `bg`. The `Ctrl+Z` keystroke will suspend (pause) the foreground process and return control to the terminal. The `bg` keystroke will resume (continue) the suspended process in the background and detach it from the terminal. The statement B is correct.

The statements A, C, and D are incorrect because they do not perform the desired task. The `bg` keystroke alone will not work unless there is a suspended process to resume. The `Ctrl+B` keystroke will not suspend the foreground process, but rather move one character backward in some applications. The `jobs` keystroke will list all processes associated with the current terminal. The `bg &` keystroke will cause an error because `bg` does not take any arguments. References: [How to Run Linux Processes in Background]

NEW QUESTION 144

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


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

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

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

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

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

Answer: AD

Explanation:

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

NEW QUESTION 149

A senior Linux administrator has created several scripts that will be used to install common system applications. These scripts are published to a repository to share with the systems team. A junior Linux administrator needs to retrieve the scripts and make them available on a local workstation. Which of the following Git commands should the junior Linux administrator use to accomplish this task?

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

Answer: C

Explanation:

To retrieve the scripts from a repository and make them available on a local workstation, the junior Linux administrator can use the command `git clone` ©. This will create a copy of the repository on the local machine, including all the scripts and history. The other commands will not clone the repository, but either fetch, checkout, or branch from an existing repository. References:

? [CompTIA Linux+ Study Guide], Chapter 10: Working with Git, Section: Cloning Repositories with Git

? [How to Clone a Git Repository]

NEW QUESTION 152

An administrator added the port 2222 for the SSH server on myhost and restarted the SSH server. The administrator noticed issues during the startup of the service. Given the following outputs:

```
$ ssh -p 2222 myhost
ssh:connect to host myhost on port 2222: Connection refused

$ nmap -p 2222 myhost
Starting Nmap 7.70 ( https://nmap.org ) at 2022-10-17 21:12 EEST
Nmap scan report for myhost (10.7.3.26)
Host is up (0.00027s latency).
rDNS record for 10.7.3.26: myhost
PORT      STATE SERVICE
2222/tcp  closed EtherNetIP-1
MAC Address: 52:54:00:F5:DF:F8 (QEMU virtual NIC)
Nmap done: 1 IP address (1 host up) scanned in 0.57 seconds

$ systemctl status sshd
● sshd.service - OpenSSH server daemon
   Loaded: loaded (/usr/lib/systemd/system/sshd.service; enabled; vendor preset: enabled)
   Active: active (running) since Mon 2022-10-17 19:40:07 CEST; 36min ago
     Docs: man:sshd(8)
           man:sshd_config(5)
  Main PID: 13186 (sshd)
    Tasks: 1 (limit: 12373)
   Memory: 1.1M
   CGroup: /system.slice/sshd.service
           └─13186 /usr/sbin/sshd -D -oCiphers=aes256-gcm@openssh.com

Oct 17 19:40:07 myhost systemd[1]: Starting OpenSSH server daemon...
Oct 17 19:40:07 myhost sshd[13186]: error: Bind to port 2222 on 0.0.0.0 failed: Permission denied.
Oct 17 19:40:07 myhost systemd[1]: Started OpenSSH server daemon.
Oct 17 19:40:07 myhost sshd[13186]: Server listening on 0.0.0.0 port 22.
```

Which of the following commands will fix the issue?

- A. `semanage port -a -t ssh_port_t -p tcp 2222`
- B. `chcon system_u:object_r:ssh_home_t /etc/ssh/*`
- C. `iptables -A INPUT -p tcp -- dport 2222 -j ACCEPT`
- D. `firewall-cmd -- zone=public -- add-port=2222/tcp`

Answer: A

Explanation:

The correct answer is A. `semanage port -a -t ssh_port_t -p tcp 2222`

This command will allow the SSH server to bind to port 2222 by adding it to the SELinux policy. The `semanage` command is a utility for managing SELinux policies. The `port` subcommand is used to manage network port definitions. The `-a` option is used to add a new record, the `-t` option is used to specify the SELinux type, the `-p` option is used to specify the protocol, and the `tcp 2222` argument is used to specify the port number. The `ssh_port_t` type is the default type for SSH ports in SELinux.

The other options are incorrect because:

* B. `chcon system_u:object_r:ssh_home_t /etc/ssh/*`

This command will change the SELinux context of all files under `/etc/ssh/` to `system_u:object_r:ssh_home_t`, which is not correct. The `ssh_home_t` type is used for user home directories that are accessed by SSH, not for SSH configuration files. The correct type for SSH configuration files is `sshd_config_t`.

* C. `iptables -A INPUT -p tcp --dport 2222 -j ACCEPT`

This command will add a rule to the iptables firewall to accept incoming TCP connections on port 2222. However, this is not enough to fix the issue, as SELinux will still block the SSH server from binding to that port. Moreover, iptables may not be the default firewall service on some Linux distributions, such as Fedora or CentOS, which use `firewalld` instead.

* D. `firewall-cmd --zone=public --add-port=2222/tcp`

This command will add a rule to the `firewalld` firewall to allow incoming TCP connections on port 2222 in the public zone. However, this is not enough to fix the issue, as SELinux will still block the SSH server from binding to that port. Moreover, `firewalld` may not be installed or enabled on some Linux distributions, such as Ubuntu or Debian, which use iptables instead.

References:

? [How to configure SSH to use a non-standard port with SELinux set to enforcing](#)

? [Change SSH Port on CentOS/RHEL/Fedora With SELinux Enforcing](#)

? [How to change SSH port when SELinux policy is enabled](#)

NEW QUESTION 153

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

A. `[root@nodea ssh —i ~/ . ssh/±d rsa root@nodeb`

B. `[root@nodea scp -i . ssh/id rsa root@nodeb`

C. `[root@nodea ssh—copy-id —i .ssh/id rsa root@nodeb`

D. `[root@nodea # ssh add -c ~/ . ssh/id rsa root@nodeb`

E. `[root@nodea # ssh add -c ~/. ssh/id rsa root@nodeb`

Answer: C

Explanation:

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

NEW QUESTION 157

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

A. Server clustering

B. Load balancing

C. RAID

D. VDI

Answer: C

Explanation:

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

NEW QUESTION 160

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

A. `tail -v 20`

B. `tail -n 20`

C. `tail -c 20`

D. `tail -l 20`

Answer: B

Explanation:

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

NEW QUESTION 165

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 169

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 174

A systems administrator needs to verify whether the built container has the app.go file in its root directory. Which of the following can the administrator use to verify the root directory has this file?

- A. docker image inspect
- B. docker container inspect
- C. docker exec <container_name> ls
- D. docker ps <container_name>

Answer: C

Explanation:

The docker exec <container_name> ls command can be used to verify whether the built container has the app.go file in its root directory. This command will run the ls command inside the specified container and list the files and directories in its root directory. If the app.go file is present, it will be displayed in the output. The docker image inspect command will display information about an image, not a container, and it will not list the files inside the image. The docker container inspect command will display information about a container, not its files. The docker ps <container_name> command is invalid, as ps does not accept a container name as an argument. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Virtualization and Cloud Technologies, page 499.

NEW QUESTION 177

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

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

Answer: B

Explanation:

The command that the administrator should run to ensure that both Java 7 and Java 8 are locally available for developers to use when deploying containers is docker image pull java:7. This command will use the docker image pull subcommand to download the java:7 image from Docker Hub, which is the default registry for Docker images. The java:7 image contains Java 7 installed on a Debian-based Linux system. The administrator can also specify a different registry by using the syntax registry/repository:tag.

The other options are not correct commands for ensuring that both Java 7 and Java 8 are locally available for developers to use when deploying containers. The docker image load java:7 command will load an image from a tar archive or STDIN, not from a registry. The docker image import java:7 command will create a new filesystem image from the contents of a tarball, not from a registry. The docker image build java:7 command will build an image from a Dockerfile, not from a registry. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks; docker image pull | Docker Docs

NEW QUESTION 180

A systems administrator is tasked with creating an Ansible playbook to automate the installation of patches on several Linux systems. In which of the following languages should the playbook be written?

- A. SQL
- B. YAML
- C. HTML
- D. JSON

Answer: B

Explanation:

The language that the playbook should be written in is YAML. YAML stands for YAML Ain't Markup Language, which is a human-readable data serialization language. YAML is commonly used for configuration files and data exchange. YAML uses indentation, colons, dashes, and brackets to represent the structure and values of the data. YAML also supports comments, variables, expressions, and functions. Ansible is an open-source tool for automating tasks and managing configuration on Linux systems. Ansible uses YAML to write playbooks, which are files that define the desired state and actions for the systems. Playbooks can be used to automate the installation of patches on several Linux systems by specifying the hosts, tasks, modules, and parameters. The language that the playbook should be written in is YAML. This is the correct answer to the question. The other options are incorrect because they are not the languages that Ansible uses for playbooks (SQL, HTML, or JSON). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 549.

NEW QUESTION 181

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

```
ls -l startup file
```

The following output is returned

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

Which of the following is MOST likely the issue?

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

Answer: A

Explanation:

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

NEW QUESTION 183

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

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

Answer: C

Explanation:

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

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

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

The other commands are incorrect because:

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

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

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

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

References:

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

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

NEW QUESTION 188

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

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

D. `systemctl start sshd`

Answer: C

Explanation:

The `systemctl reload sshd` command can be used to apply the configuration changes of the SSH server daemon without restarting it. This is useful to avoid interrupting existing connections. The `systemctl stop sshd` command would stop the SSH server daemon, not apply the changes. The `systemctl mask sshd` command would prevent the SSH server daemon from being started, not apply the changes. The `systemctl start sshd` command would start the SSH server daemon if it is not running, but it would not apply the changes if it is already running. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Secure Shell (SSH), page 415.

NEW QUESTION 191

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
<code>/dev/sdb1</code>	100G	50G	50G	50%	<code>/oracle1</code>

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 193

An application developer received a file with the following content:

```
##This is a sample Image ## FROM ubuntu:18.04
MAINTAINER demohut@gmail.com.hac COPY ./app
RUN make /app
CMD python /app/app.py RUN apt-get update
RUN apt-get install -y nginx CMD ["echo","Image created"]
```

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

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

Answer: A

Explanation:

The `docker build` command is used to build an image from a Dockerfile and a context¹. The Dockerfile is a text file that contains the instructions for creating the image, and the context is a set of files that can be used in the image creation process¹. The file that the developer received is an example of a Dockerfile.

The `-t` option is used to specify a name and an optional tag for the image¹. The name and tag are separated by a colon (:), and the tag is usually used to indicate the version of the image². For example, `-t myimage:1.0` means that the image will be named `myimage` and tagged as `1.0`.

The last argument of the `docker build` command is the path to the context, which can be a local directory or a URL¹. The dot (.) means that the current working directory is the context². Therefore, `docker build -t myimage:1.0 .` means that the image will be built from the Dockerfile and the files in the current working directory, and it will be named `myimage` and tagged as `1.0`.

NEW QUESTION 195

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

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

Answer: A

Explanation:

To create a symlink for `/usr/local/bin/app-a`, which was installed in `/usr/local/share/app-a`, the administrator can use the command `ln -s /usr/local/share/app-a /usr/local/bin/app-a` (A). This will create a symbolic link named `/usr/local/bin/app-a` that points to the original file `/usr/local/share/app-a`. The other commands will not create a symlink, but either move, copy, or synchronize the file. References:

? [CompTIA Linux+ Study Guide], Chapter 3: Working with Files, Section: Creating Links

? [How to Create Symbolic Links in Linux]

NEW QUESTION 197

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

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

Answer: B

Explanation:

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

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

? [How to Use Fdisk Command in Linux]

NEW QUESTION 199

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

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

Answer: C

Explanation:

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

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

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

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

NEW QUESTION 204

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:$6$S3uOw6qWx9876jGhgKJedfH987634534voj.:18883:0:99999:7:::
```

Which of the following commands 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:

The command pam_tally2 -u joe -r will resolve the issue of Joe being unable to log in to the Linux system. The pam_tally2 command is a tool for managing the login counter for the PAM (Pluggable Authentication Modules) system. PAM is a framework for managing authentication and authorization on Linux systems. PAM allows the administrator to define the rules and policies for accessing various system resources and services, such as login, sudo, ssh, or cron. PAM also supports different types of authentication methods, such as passwords, tokens, biometrics, or smart cards. PAM can be used to implement login restrictions, such as limiting the number of failed login attempts, locking the account after a certain number of failures, or enforcing a minimum or maximum time between login attempts. The pam_tally2 command can display, reset, or unlock the login counter for the users or hosts. The -u joe option specifies the user name that the command should apply to. The -r option resets the login counter for the user. The command pam_tally2 -u joe -r will reset the login counter for Joe, which will unlock his account and allow him to log in to the Linux system. This will resolve the issue of Joe being unable to log in to the Linux system. This is the correct command to use to resolve the issue. The other options are incorrect because they either do not unlock the account (usermod -s /bin/bash joe or passwd -u joe) or do not affect the login counter (chage -E 90 joe).

References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Implementing Basic Security, page 517.

NEW QUESTION 205

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

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

Answer: C

Explanation:

The issue is that the administrator did not set the default target to graphical.target. A target is a unit of systemd that groups together other units by a common purpose or state. The graphical.target is a target that starts the graphical user interface (GUI) along with other services. The administrator used the command `systemctl isolate graphical.target` to switch to this target temporarily, but this does not change the default target that is activated at boot time. To make this change permanent, the administrator should have used the command `systemctl set-default graphical.target`, which creates a symbolic link from `/etc/systemd/system/default.target` to `/usr/lib/systemd/system/graphical.target`.

The other options are not correct explanations for the issue. The administrator did reboot the server properly by using `systemctl reboot`, which shuts down and restarts the system cleanly. The administrator did not need to set the default target to basic.target, which is a minimal target that only starts essential services. The administrator did not shut down the server improperly, which could have caused file system corruption or data loss, but not affect the default target. References: `systemctl(1)` - Linux manual page; How to Change Runlevels (targets) in SystemD

NEW QUESTION 206

A systems administrator is working on a security report from the Linux servers. Which of the following commands can the administrator use to display all the firewall rules applied to the Linux servers? (Select two).

- A. `ufw limit`
- B. `iptables -F`
- C. `systemctl status firewalld`
- D. `firewall-cmd --list-all`
- E. `ufw status`
- F. `iptables -A`

Answer: DE

Explanation:

These commands can display all the firewall rules applied to the Linux servers, depending on which firewall service is being used.

? The `firewall-cmd` command is a utility for managing `firewalld`, which is a dynamic firewall service that supports zones and services. The `--list-all` option will show all the settings and rules for the default zone, or for a specific zone if specified. For example, `firewall-cmd --list-all --zone=public` will show the rules for the public zone¹.

? The `ufw` command is a frontend for `iptables`, which is a low-level tool for manipulating `netfilter`, the Linux kernel's packet filtering framework. The `status` option will show the status of `ufw` and the active rules, or the numbered rules if `verbose` is specified. For example, `ufw status verbose` will show the numbered rules and other information².

The other options are incorrect because:

* A. `ufw limit`

This command will limit the connection attempts to a service or port using `iptables`' recent module. It does not display any firewall rules².

* B. `iptables -F`

This command will flush (delete) all the rules in the selected chain, or all chains if none is given. It does not display any firewall rules³.

* C. `systemctl status firewalld`

This command will show the status of the `firewalld` service, including whether it is active or not, but it does not show the firewall rules⁴.

* F. `iptables -A`

This command will append one or more rules to the end of the selected chain. It does not display any firewall rules³.

NEW QUESTION 210

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

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

Answer: A

Explanation:

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

NEW QUESTION 213

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

```
# ip -s link list dev enp0s25

2: enp0s25: <BROADCAST,MULTICAST,LOWER_UP> mtu 1500 qdisc fq_codel state DOWN mode DEFAULT group default qlen 1000    link/ether
ac:12:34:56:78:cd brd ff:ff:ff:ff:ff:ff

    RX: bytes    packets  errors  dropped missed  mcast
       2011664755 3579033   2394390 508      0        0

    TX: bytes    packets  errors  dropped carrier collsns
       309541780 1705408    0        0    12340     0
```

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

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

Answer: B

Explanation:

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

? CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Networking, Section: Troubleshooting Network Issues, Page 359.
? Linux+ (Plus) Certification, Exam Objectives: 4.3 Given a scenario, troubleshoot and resolve basic network configuration and connectivity issues.

NEW QUESTION 216

A user is attempting to log in to a Linux server that has Kerberos SSO enabled. Which of the following commands should the user run to authenticate and then show the ticket grants? (Select TWO).

- A. kinit
- B. klist
- C. kexec
- D. kload
- E. pkexec
- F. realm

Answer: AB

Explanation:

The following commands can help the user to authenticate and show the ticket grants using Kerberos SSO on a Linux server:

? kinit: This command obtains and caches an initial ticket-granting ticket (TGT) for

the user from the Kerberos key distribution center (KDC). The user needs to enter their password or use a keytab file to authenticate1.

? klist: This command lists the cached tickets, including the TGT and any service tickets, for the user. It also shows the expiration time and flags for each ticket2.

For example, the user can run the following commands to log in and view their tickets:

```
$ kinit username@REALM Password for username@REALM:
```

```
$ klist
```

```
Ticket cache: FILE:/tmp/krb5cc_1000 Default principal: username@REALM
```

```
Valid starting Expires Service principal
```

```
04/06/2023 16:06:59 04/07/2023 02:06:59 krbtgt/REALM@REALM
```

```
renew until 04/13/2023 16:06:59 References:
```

? kinit(1) - Linux man page, section "Description".

? klist(1) - Linux man page, section "Description".

NEW QUESTION 219

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

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

Answer: A

Explanation:

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

NEW QUESTION 222

A Linux administrator provisioned a new web server with custom administrative permissions for certain users. The administrator receives a report that user1 is unable to restart the Apache web service on this server. The administrator reviews the following output:

```
[ root@server ] # id user1
```

```
UID=1011 (user1) gid=1011 (USER1) groups=1011 (user1), 101 (www-data), 1120 (webadmin)
```

```
[ root@server ] # cat /etc/sudoers.d/custom.conf
```

```
user1 ALL=/usr/sbin/systemctl start httpd, /usr/sbin/systemctl stop httpd webadmin ALL=NOPASSWD: /etc/init.d.httpd restart, /sbin/service httpd restart,
```

```
/usr/sbin/apache2ctl restart
```

```
#!/usr/bin/perl ALL=(ALL) NOPASSWD: ALL
```

Which of the following would most likely resolve the issue while maintaining a least privilege security model?

- A. User1 should be added to the wheel group to manage the service.
- B. User1 should have "NOPASSWD:" after the "ALL=" in the custo
- C. conf.
- D. The wheel line in the custo
- E. conf file should be uncommented.
- F. Webadmin should be listed as a group in the custo
- G. conf file.

Answer: D

Explanation:

The custom.conf file grants sudo privileges to user1 and webadmin for managing the Apache web service, but it uses different commands for each of them. User1 is allowed to use systemctl to start and stop the httpd service, while webadmin is allowed to use init.d, service, or apache2ctl to restart the httpd service. However, the user1 is unable to restart the service, only start and stop it. To fix this, user1 should be able to use the same commands as webadmin, which can be achieved by listing webadmin as a group in the custom.conf file, using the syntax %groupname. This way, user1 will inherit the sudo privileges of the webadmin group, and be able to restart the Apache web service without compromising the least privilege security model.

References

? Sudo and Sudoers Configuration | Servers for Hackers, section "Groups"

? Chapter 12. Managing sudo access - Red Hat Customer Portal, section "12.1."

Configuring sudo access for users and groups”

NEW QUESTION 224

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

XK0-005 Practice Exam Features:

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

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