

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

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

Answer: A

Explanation:

This command will load the KVM module as well as any related dependencies, such as kvm-intel or kvm-amd, depending on the processor type. The modprobe command is a Linux utility that reads the /etc/modules.conf file and adds or removes modules from the kernel. It also resolves any dependencies between modules, so that they are loaded in the correct order.
 The other options are incorrect because:
 * B. insmod kvm
 This command will only load the KVM module, but not any related dependencies. The insmod command is a low-level Linux utility that inserts a single module into the kernel. It does not resolve any dependencies between modules, so they have to be loaded manually.
 * C. depmod kvm
 This command will not load the KVM module at all, but only create a list of module dependencies for modprobe to use. The depmod command is a Linux utility that scans the installed modules and generates a file called modules.dep that contains dependency information for each module.
 * D. hotplug kvm
 This command is invalid and does not exist. The hotplug mechanism is a feature of the Linux kernel that allows devices to be added or removed while the system is running. It does not have anything to do with loading modules.

NEW QUESTION 3

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 4

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

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

Answer: D

Explanation:

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

NEW QUESTION 5

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

```
$ sudo visudo
```

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

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

Which of the following actions will resolve this issue?

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

Answer: B

NEW QUESTION 6

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

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

Answer: B

Explanation:

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

NEW QUESTION 7

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 8

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

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

Answer: A

Explanation:

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

NEW QUESTION 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. rmmod
- B. ls -ll /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 information¹.

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

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

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

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

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

NEW QUESTION 10

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

Hostname: devel.comptia.org

IP address: 5.5.5.1, 5.5.5.2, 5.5.5.3, 5.5.5.4

Name server: 5.5.5.254

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

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

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

Answer: BDE

Explanation:

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

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

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

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

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

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

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

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

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

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

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

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

NEW QUESTION 10

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

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

Answer: A

Explanation:

The best command sequence to rebuild a kernel module from source code is A. ./configure make make install. This is the standard way to compile and install a Linux kernel module, as explained in the web search result ⁵. The other commands are either not relevant, not valid, or not sufficient for this task. For example:

? B. wget gcc cp will try to download, compile, and copy a file, but it does not specify the source code, the module name, or the destination directory.

? C. tar xvzf build cp will try to extract, build, and copy a compressed file, but it does not specify the file name, the module name, or the destination directory.

? D. build install configure will try to run three commands that are not defined or recognized by the Linux shell.

NEW QUESTION 12

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 13

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

NEW QUESTION 16

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

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

Answer: B

Explanation:

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

NEW QUESTION 20

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 25

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 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 32

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 34

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

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

Answer: C

Explanation:

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

NEW QUESTION 36

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

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

Answer: C

Explanation:

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

NEW QUESTION 40

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

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

Answer: A

Explanation:

The best command to use to prevent the `httpd` web service from being started both manually and automatically on a server is `A. systemctl mask httpd`. This command will create a symbolic link from the `httpd` service unit file to `/dev/null`, which will make the service impossible to start or enable. This is different from

systemctl disable httpd, which will only prevent the service from starting automatically on boot, but not manually. The other commands are either not relevant or not sufficient for this task. For example:

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

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

NEW QUESTION 41

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 47

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

Which of the following is correct based on the output received from the executed command?

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

Answer: C

Explanation:

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

The other options are not correct based on the output received from the executed command. The server's CPU is not taking too long to process users' requests, because the us state, which stands for user, is only 0.2%, which means that the CPU is barely used by user processes. The server's CPU does not show a high idle-time value, because the id state, which stands for idle, is 0.0%, which means that the CPU is not idle at all. The server's CPU value for the time spent on system processes is not low, because the sy state, which stands for system, is 33.1%, which means that the CPU is heavily used by system processes.

References: How to Use the Linux top Command (and Understand Its Output); [Understanding Linux CPU Load - when should you be worried?]

NEW QUESTION 50

A User on a Linux workstation needs to remotely start an application on a Linux server and then forward the graphical display of that application back to the Linux workstation. Which of the following would enable the user to perform this action?

- A. ssh -X user@server application
- B. ssh -y user@server application
- C. ssh user@server application
- D. ssh -D user@server application

Answer: A

Explanation:

The ssh -X option enables X11 forwarding, which allows the user to run graphical applications on the remote server and display them on the local workstation. The user needs to specify the username, the server address, and the application name after the ssh -X command. The remote server also needs to have X11Forwarding enabled and xauth installed for this to work. References:

? The web search result 8 explains how to run a GUI application through SSH by configuring both the SSH client and server.

? The web search result 6 provides a detailed answer on how to forward X over SSH to run graphics applications remotely, with examples and troubleshooting tips.

? The CompTIA Linux+ Certification Exam Objectives mention that the candidate should be able to “use SSH for remote access and management” as part of the System Operation and Maintenance domain¹.

NEW QUESTION 52

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 56

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

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

Answer: D

Explanation:

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

NEW QUESTION 60

A Linux engineer needs to block an incoming connection from the IP address 2.2.2.2 to a secure shell server and ensure the originating IP address receives a response that a firewall is blocking the connection. Which of the following commands can be used to accomplish this task?

- A. iptables -A INPUT -p tcp -- dport ssh -s 2.2.2.2 -j DROP
- B. iptables -A INPUT -p tcp -- dport ssh -s 2.2.2.2 -j RETURN
- C. iptables -A INPUT -p tcp -- dport ssh -s 2.2.2.2 -j REJECT
- D. iptables -A INPUT -p tcp -- dport ssh -s 2.2.2.2 -j QUEUE

Answer: C

Explanation:

The REJECT target sends back an error packet to the source IP address, indicating that the connection is refused by the firewall. This is different from the DROP target, which silently discards the packet without any response. The RETURN target returns to the previous chain, which may or may not accept the connection. The QUEUE target passes the packet to a userspace application for further processing, which is not the desired outcome in this case.

References

? CompTIA Linux+ (XK0-005) Certification Study Guide, page 316

? iptables - ssh - access from specific ip only - Server Fault, answer by Eugene Ionichev

NEW QUESTION 65

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

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

Answer: D

Explanation:

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

To insert a character at the end of the current line in the vi editor, the user can press the A key in command mode. This will move the cursor to the end of the line and switch to insert mode. Then, the user can type the desired character and press Esc to return to command mode. The statement D is correct. The statements A, B, C, and E are incorrect because they do not perform the desired task. The p key in command mode will paste the previously copied or deleted text after the cursor. The r key in command mode will replace the character under the cursor with another character. The bb key in command mode will move the cursor back two words. The i key in command mode will switch to insert mode before the cursor. References: [How to Use vi Text Editor in Linux]

NEW QUESTION 66

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

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

Answer: A

Explanation:

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

NEW QUESTION 71

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 73

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 75

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

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

Answer: ACE

Explanation:

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

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

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

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

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

The lvresize command is a tool for resizing logical volumes on Linux systems. The -L option specifies the new size of the logical volume, in this case +10G, which means 10 GB more than the current size. The -r option resizes the underlying file system as well. The /dev/mapper/rootvg-rootlv is the device name of the logical volume, which is the same as the partition name. The command lvresize -L +10G -r /dev/mapper/rootvg-rootlv will increase the size of the logical volume and the root filesystem by 10 GB and free up space. This command will help the administrator solve the problem and free up space.

The other options are incorrect because they either do not affect the root filesystem (fdisk -l /dev/sdb, pvcreate /dev/sdb, lsblk /dev/sda, or vgextend /dev/rootvg /dev/sdb) or do not use the correct syntax (fdisk -l /dev/sdb instead of fdisk -l /dev/sdb or parted -l /dev/mapper/rootvg-rootlv instead of parted /dev/mapper/rootvg-rootlv print). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 318-319, 331-332.

NEW QUESTION 78

A developer needs to launch an Nginx image container, name it Web001, and ex-pose 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 Web00l 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 83

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 86

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

Path not found

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

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

Answer: B

Explanation:

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

NEW QUESTION 89

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

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

Answer: BF

Explanation:

"modify the /etc/fstab text file to automatically mount the new partition by opening it in an editor and adding the following line:

/dev/ xxx 1 /data ext4 defaults 1 2

where xxx is the device name of the storage device"

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

NEW QUESTION 94

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

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

Answer: A

Explanation:

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

References

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

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

NEW QUESTION 96

A Linux administrator copied a Git repository locally, created a feature branch, and committed some changes to the feature branch. Which of the following Git actions should the Linux administrator use to publish the changes to the main branch of the remote repository?

- A. rebase
- B. tag
- C. commit
- D. push

Answer: D

Explanation:

The push action is used to publish the changes made in a local branch to a remote branch of a Git repository. This action will update the remote branch with the commits made in the local branch and synchronize the two branches. The rebase action is used to reapply commits from one branch onto another branch, creating a linear history of commits. This action does not publish any changes to a remote repository. The tag action is used to create an annotated reference to a specific commit in a Git repository. This action does not publish any changes to a remote repository. The commit action is used to record changes made in the local repository and create a new snapshot of the project state. This action does not publish any changes to a remote repository. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 20: Writing and Executing Bash Shell Scripts, page 579.

NEW QUESTION 97

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

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

Answer: C

Explanation:

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

NEW QUESTION 100

A Linux administrator wants to set the SUID of a file named dev_team.txt with 744 access rights. Which of the following commands will achieve this goal?

- A. `chmod 4744 dev_team.txt`
- B. `chmod 744 --setuid dev_team.txt`
- C. `chmod -c 744 dev_team.txt`
- D. `chmod -v 4744 --suid dev_team.txt`

Answer: A

Explanation:

The command that will set the SUID of a file named dev_team.txt with 744 access rights is `chmod 4744 dev_team.txt`. This command will use the `chmod` utility to change the file mode bits of dev_team.txt. The first digit (4) represents the SUID bit, which means that when someone executes dev_team.txt, it will run with the permissions of the file owner. The next three digits (744) represent the read, write, and execute permissions for the owner (7), group (4), and others (4). This means that the owner can read, write, and execute dev_team.txt, while the group and others can only read it.

The other options are not correct commands for setting the SUID of a file with 744 access rights. The `chmod 744 --setuid dev_team.txt` command is invalid because there is no `--setuid` option in `chmod`. The `chmod -c 744 dev_team.txt` command will change the file mode bits to 744, but it will not set the SUID bit. The `-c` option only means that `chmod` will report when a change is made. The `chmod -v 4744 --suid dev_team.txt` command is also invalid because there is no `--suid` option in `chmod`. The `-v` option only means that `chmod` will output a diagnostic for every file processed. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing Users and Groups; `chmod(1)` - Linux manual page

NEW QUESTION 105

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 110

An administrator needs to make some changes in the IaC declaration templates. Which of the following commands would maintain version control?

- A. `git clone https://github.com/comptia/linux+- .git git push origin`

- B. git clone https://qithub.com/comptia/linux+- .git git fetch New-Branch
- C. git clone https://github.com/comptia/linux+- .git git status
- D. git clone https://github.com/comptia/linux+- .git git checkout -b <new-branch>

Answer: D

Explanation:

The command that will maintain version control while making some changes in the IaC declaration templates is git checkout -b <new-branch>. This command uses the git tool, which is a distributed version control system that tracks changes in source code and enables collaboration among developers. The checkout option switches to a different branch in the git repository, where a branch is a pointer to a specific commit in the history. The -b option creates a new branch with the given name, and switches to it. This way, the administrator can make changes in the new branch without affecting the main branch, and later merge them if needed.

The other options are not correct commands for maintaining version control while making some changes in the IaC declaration templates. The git clone https://github.com/comptia/linux+- .git command will clone an existing repository from a remote URL to a local directory, but it will not create a new branch for making changes. The git push origin command will push the local changes to a remote repository named origin, but it will not create a new branch for making changes. The git fetch New-Branch command will fetch updates from a remote branch named New-Branch, but it will not create a new branch for making changes. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Source Code; Git - Basic Branching and Merging

NEW QUESTION 112

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

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

Answer: B

Explanation:

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

NEW QUESTION 117

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 120

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 124

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

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

Answer: C

Explanation:

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

NEW QUESTION 125

An administrator is trying to diagnose a performance issue and is reviewing the following output:

```
avg-cpu:  %user   %nice   %system %iowait  %steal   %idle
           2.00    0.00    3.00    32.00    0.00    63.00

Device            tps    kB_read/s    kB_wrtn/s    kB_read    kB_wrtn
sdb                345.00         0.02         0.04 4739073123 23849523
sdb1               345.00    32102.03    12203.01 4739073123 23849523
```

System Properties: CPU: 4 vCPU

Memory: 40GB

Disk maximum IOPS: 690

Disk maximum throughput: 44Mbps | 44000Kbps

Based on the above output, which of the following BEST describes the root cause?

- A. The system has reached its maximum IOPS, causing the system to be slow.
- B. The system has reached its maximum permitted throughput, therefore `iowait` is increasing.
- C. The system is mostly idle, therefore the `iowait` is high.
- D. The system has a partitioned disk, which causes the IOPS to be doubled.

Answer: B

Explanation:

The system has reached its maximum permitted throughput, therefore `iowait` is increasing. The output of `iostat -x` shows that the device `sda` has an average throughput of 44.01 MB/s, which is equal to the disk maximum throughput of 44 Mbps. The output also shows that the device `sda` has an average `iowait` of 99.99%, which means that the CPU is waiting for the disk to complete the I/O requests. This indicates that the disk is the bottleneck and the system is slow due to the high `iowait`. The other options are incorrect because they are not supported by the outputs. The system has not reached its maximum IOPS, as the device `sda` has an average IOPS of 563.50, which is lower than the disk maximum IOPS of 690. The system is not mostly idle, as the output of `top` shows that the CPU is 100% busy. The system does not have a partitioned disk, as the output of `lsblk` shows that the device `sda` has only one partition `sda1`. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Optimizing Linux Systems, pages 513-514.

NEW QUESTION 126

An administrator would like to securely connect to a server and forward port 8080 on a local machine to port 80 on the server. Which of the following commands should the administrator use to satisfy both requirements?

- A. `ssh -L 8080:localhost:80 admin@server`
- B. `ssh -R 8080:localhost:80 admin@server`
- C. `ssh -L 80 :localhost:8080 admin@server`
- D. `ssh -R 80 :localhost:8080 admin@server`

Answer: A

Explanation:

This command will create a local port forwarding, which means that connections from the SSH client are forwarded via the SSH server, then to a destination server. In this case, the destination server is the same as the SSH server (localhost), and the destination port is 80. The SSH client will listen on port 8080 on the

local machine, and any connection to that port will be forwarded to port 80 on the server. This way, the administrator can securely access the web service running on port 80 on the server by using `http://localhost:8080` on the local machine.

The other options are incorrect because:

* B. `ssh -R 8080:localhost:80 admin@server`

This command will create a remote port forwarding, which means that connections from the SSH server are forwarded via the SSH client, then to a destination server. In this case, the destination server is the same as the SSH client (localhost), and the destination port is 80. The SSH server will listen on port 8080 on the remote machine, and any connection to that port will be forwarded to port 80 on the client. This is not what the administrator wants to do.

* C. `ssh -L 80:localhost:8080 admin@server`

This command will also create a local port forwarding, but it will use port 80 on the local machine and port 8080 on the server. This is not what the administrator wants to do, and it may also fail if port 80 is already in use by another service on the local machine.

* D. `ssh -R admin@server`

This command is incomplete and invalid. It does not specify any port numbers or destination addresses for the remote port forwarding. It will also fail if the SSH server does not allow remote port forwarding.

References:

? CompTIA Linux+ Certification Exam Objectives

? How to Set up SSH Tunneling (Port Forwarding)

NEW QUESTION 130

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

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

Answer: A

Explanation:

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

NEW QUESTION 133

Due to performance issues on a server, a Linux administrator needs to terminate an unresponsive process. Which of the following commands should the administrator use to terminate the process immediately without waiting for a graceful shutdown?

- A. `kill -SIGKILL 5545`
- B. `kill -SIGTERM 5545`
- C. `kill -SIGHUP 5545`
- D. `kill -SIGINT 5545`

Answer: A

Explanation:

To terminate an unresponsive process immediately without waiting for a graceful shutdown, the administrator can use the command `kill -SIGKILL 5545` (A). This will send a signal to the process with the PID 5545 that cannot be ignored or handled by the process, and force it to stop. The other commands will send different signals that may allow the process to perform some cleanup or termination actions, or may be ignored by the process. References:

? [CompTIA Linux+ Study Guide], Chapter 6: Managing Processes, Section: Killing Processes

? [How to Kill Processes in Linux]

NEW QUESTION 138

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 139

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 143

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 148

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 153

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

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

Answer: A

Explanation:

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

NEW QUESTION 158

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 163

A new Linux systems administrator just generated a pair of SSH keys that should allow connection to the servers. Which of the following commands can be used to copy a key file to remote servers? (Choose two.)

- A. `wget`
- B. `ssh-keygen`
- C. `ssh-keyscan`
- D. `ssh-copy-id`
- E. `ftpd`
- F. `scp`

Answer: DF

Explanation:

The commands `ssh-copy-id` and `scp` can be used to copy a key file to remote servers. The command `ssh-copy-id` copies the public key to the `authorized_keys` file

on the remote server, which allows the user to log in without a password. The command scp copies files securely over SSH, which can be used to transfer the key file to any location on the remote server. The other options are incorrect because they are not related to copying key files. The command wget downloads files from the web, the command ssh-keygen generates key pairs, the command ssh-keyscan collects public keys from remote hosts, and the command ftpd is a FTP server daemon. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, pages 408-410.

NEW QUESTION 168

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 172

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

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

Answer: D

Explanation:

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

NEW QUESTION 177

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

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

Answer: C

Explanation:

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

NEW QUESTION 179

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

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

Answer: A

Explanation:

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

sudo passwd tom

The other options are incorrect because:

* B. sudo userdel

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

* C. sudo chage

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

* D. sudo usermod

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

References:

? How to Change Account Passwords on Linux

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

? CompTIA Linux+ Certification Exam Objectives

NEW QUESTION 180

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

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

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 181

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

Output 1:

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

Output 2:

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 186

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

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

Answer: B

Explanation:

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

NEW QUESTION 190

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

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

Answer: A

Explanation:

The best command to use to list all current containers, regardless of their running state, is A. `docker ps -a`. This command will show all containers, both running and stopped, with details such as container ID, image name, status, and ports. The other commands are either invalid or not relevant for this task. For example:
? B. `docker list` is not a valid command. There is no subcommand named `list` in `docker`.
? C. `docker image ls` will list all the images available on the local system, not the containers.
? D. `docker inspect image` will show detailed information about a specific image, not all the containers.

NEW QUESTION 191

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

- A. `unzip -v`
- B. `bzip2 -z`
- C. `gzip`
- D. `funzip`

Answer: C

Explanation:

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

NEW QUESTION 193

A Linux system is having issues. Given the following outputs:

```
# dig @192.168.2.2 mycomptiahost
; << >> DiG 9.9.4-RedHat-9.9.4-74.el7_6.1 << >> @192.168.2.2 mycomptiahost
; (1 server found)
;; global options: +cmd
;; connection timed out; no servers could be reached
# nc -v 192.168.2.2 53
Ncat: Version 7.70 ( https://nmap.org/ncat ) Ncat: Connection timed out.
# ping 192.168.2.2
PING 192.168.2.2 (192.168.2.2) 56(84) bytes of data.
64 bytes from 192.168.2.2: icmp_seq=1 ttl=117 time=4.94 ms 64 bytes from 192.168.2.2: icmp_seq=2 ttl=117 time=10.5 ms
```

Which of the following best describes this issue?

- A. The DNS host is down.
- B. The name `mycomptiahost` does not exist in the DNS.
- C. The Linux engineer is using the wrong DNS port.
- D. The DNS service is currently not available or the corresponding port is blocked.

Answer: D

Explanation:

The ping command shows that the Linux system can reach the DNS server at 192.168.2.2, so the DNS host is not down. The dig and nc commands show that the Linux system cannot connect to the DNS server on port 53, which is the standard port for DNS queries. This means that either the DNS service is not running on the DNS server, or there is a firewall or network device blocking the port 53 traffic. Therefore, the DNS service is currently not available or the corresponding port is blocked. References: 1: How To Troubleshoot DNS Client Issues in Linux - RootUsers 2: 6 Best Tools to Troubleshoot DNS Issues in Linux - Tecmint 3: How To Troubleshoot DNS in Linux - OrcaCore 4: Fixing DNS Issues in Ubuntu 20.04 | DeviceTests

NEW QUESTION 196

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

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

Which of the following commands would resolve the issue?

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

Answer: D

Explanation:

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

The other options are not correct commands for resolving the issue. The `restorecon .ssh/authorized_keys` command will restore the default SELinux security context for the `.ssh/authorized_keys` file, but this will not change its permissions or ownership. The `ssh_keygen -t rsa -o .ssh/authorized_keys` command is invalid because `ssh_keygen` is not a valid command (the correct command is `ssh-keygen`), and the `-o` option is used to specify a new output format for the key file, not the output file name. The `chown root:root .ssh/authorized_keys` command will change the owner and group of the `.ssh/authorized_keys` file to root, but this will not change its permissions or make it accessible by the user who wants to log in with SSH key authentication. References: How to Use Public Key Authentication with SSH; `chmod(1)` - Linux manual page

NEW QUESTION 197

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

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

Answer: D

Explanation:

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

NEW QUESTION 202

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 204

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 208

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

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

Answer: D

NEW QUESTION 211

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

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

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

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

Answer: D

Explanation:

The command `setfacl -m g:finance:rw file` will permanently fix the access issue while limiting access to IT and finance department employees. The `setfacl` command is a tool for modifying the access control lists (ACLs) of files and directories on Linux systems. The ACLs are a mechanism that allows more fine-grained control over the permissions of files and directories than the traditional owner-group-others model. The `-m` option specifies the modification to the ACL. The `g:finance:rw` means that the group named finance will have read and write permissions on the file. The file is the name of the file to modify, in this case /opt/work/file. The command `setfacl -m g:finance:rw file` will add an entry to the ACL of the file that will grant read and write access to the finance group. This will fix the access issue and allow the finance employees to access the file. The command will also preserve the existing permissions of the file, which means that the IT employees will still have read and write access to the file. This will limit the access to IT and finance department employees and prevent unauthorized access from other users.

This is the correct command to use to accomplish the task. The other options are incorrect because they either do not fix the access issue (`chattr +i file` or `chown it:finance file`) or do not limit the access to IT and finance department employees (`chmod 666 file`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing File Permissions and Ownership, page 352.

NEW QUESTION 216

A systems administrator is tasked with creating a cloud-based server with a public IP address.

```
---
-name: start an instance with a public IP address
community.abc.ec2_instance:
  name: "public-compute-instance"
  key_name: "comptia-ssh-key"
  vpc_subnet_id: subnet-5cjs1
  instance_type: instance.type
  security_group: comptia
  network:
    assign_public_ip: true
  image_id: ami-1234568
  tags:
    Environment: Comptia-Items-Writing-Workshop
...
```

Which of the following technologies did the systems administrator use to complete this task?

- A. Puppet
- B. Git
- C. Ansible
- D. Terraform

Answer: D

Explanation:

The systems administrator used Terraform to create a cloud-based server with a public IP address. Terraform is a tool for building, changing, and versioning infrastructure as code. Terraform can create and manage resources on different cloud platforms, such as AWS, Azure, or Google Cloud. Terraform uses a declarative syntax to describe the desired state of the infrastructure and applies the changes accordingly. Terraform can also assign a public IP address to a cloud server by using the appropriate resource attributes. This is the correct technology that the systems administrator used to complete the task. The other options are incorrect because they are either not designed for creating cloud servers (Puppet or Git) or not capable of assigning public IP addresses (Ansible). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 559.

NEW QUESTION 220

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

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

Answer: C

Explanation:

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

NEW QUESTION 223

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 -l`
- 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 226

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

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

Answer: A

Explanation:

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

NEW QUESTION 227

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

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

Answer: D

Explanation:

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

The other options are incorrect because:

* A. git branch -m staging

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

* B. git commit -m staging

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

* C. git status -b staging

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

? Git - git-checkout Documentation

? Git Tutorial: Create a New Branch With Git Checkout

? Git Branching - Basic Branching and Merging

NEW QUESTION 230

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 233

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

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

Answer: D

Explanation:

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

NEW QUESTION 235

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

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

Answer: CF

Explanation:

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

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

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

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

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

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

NEW QUESTION 240

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

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

Answer: C

Explanation:

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

NEW QUESTION 243

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 244

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 245

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 246

.....

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](#)