

# CompTIA

## Exam Questions XK0-005

CompTIA Linux+ Certification Exam



**NEW QUESTION 1**

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

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

**Answer: A**

**Explanation:**

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

The other options are incorrect because:

\* B. insmod kvm

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

\* C. depmod kvm

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

\* D. hotplug kvm

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

**NEW QUESTION 2**

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

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

**Answer: C**

**Explanation:**

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

The other options are not correct commands for backing up a directory to a remote server. The scp -p /data remote:/backup/data command will copy the /data directory as a file named data under /backup/ on the remote server, not as a subdirectory with its contents. The -p option preserves the permissions and timestamps of the file, but not the ownership or symbolic links. The ssh -i /remote:/backup/ /data command will try to use /remote:/backup/ as an identity file for SSH authentication, which is not valid. The cp -r

/data /remote/backup/ command will try to copy the /data directory to a local directory named /remote/backup/, not to a remote server. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks; rsync(1) - Linux manual page

**NEW QUESTION 3**

A 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 4**

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

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

**Answer: D**

**Explanation:**

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

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

References: rpm(8) - Linux manual

page; Using RPM to Verify Installed Packages

#### NEW QUESTION 5

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

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

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

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

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

```
[Install]
```

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

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

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

**Answer:** AE

#### Explanation:

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

#### NEW QUESTION 6

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

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

**Answer:** C

#### Explanation:

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

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

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

#### NEW QUESTION 7

A Linux administrator is troubleshooting a memory-related issue. Based on the output of the commands:

```
$ vmstat -s --unit M
```

```
968 M total memory
331 M used memory
482 M active memory
279 M inactive memory
99 M free memory
```

```
$ free -h
```

	total	used	free	shared	buff/cache	available
Mem:	968M	331M	95M	13M	540M	458M
Swap:	0	0	0			

```
$ ps -aux | grep script.sh
```

USER	PID	%CPU	%MEM	VSZ	RSS	TTY	STAT	START	TIME	COMMAND
user	8321	2.8	40.5	3224846	371687	7	SN	16:49	2:09	/home/user/script.sh

Which of the following commands would address the issue?

- A. top -p 8321
- B. kill -9 8321
- C. renice -10 8321
- D. free 8321

**Answer:** B

**Explanation:**

The command that would address the memory-related issue is kill -9 8321. This command will send a SIGKILL signal to the process with the PID 8321, which is the mysqld process that is using 99.7% of the available memory according to the top output. The SIGKILL signal will terminate the process immediately and free up the memory it was using. However, this command should be used with caution as it may cause data loss or corruption if the process was performing some critical operations.

The other options are not correct commands for addressing the memory-related issue. The top -p 8321 command will only display information about the process with the PID 8321, but will not kill it or reduce its memory usage. The renice -10 8321 command will change the priority (niceness) of the process with the PID 8321 to -10, which means it will have a higher scheduling priority, but this will not affect its memory consumption. The free 8321 command is invalid because free does not take a PID as an argument; free only displays information about the total, used, and free memory in the system. References: How to troubleshoot Linux server memory issues; kill(1) - Linux manual page

**NEW QUESTION 8**

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 9**

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 10**

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 10**

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

ping: comptia.org: Name or service not known

Which of the following files should the administrator verify?

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

**Answer:** C

**Explanation:**

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

```
nameserver 8.8.8.8 nameserver 8.8.4.4
```

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

**NEW QUESTION 15**

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 19**

The development team wants to prevent a file from being modified by all users in a Linux system, including the root account. Which of the following commands can be used to accomplish this objective?

- A. chmod / app/conf/file
- B. setenforce / app/ conf/ file
- C. chattr +i /app/conf/file
- D. chmod 0000 /app/conf/file

**Answer: C**

**Explanation:**

The chattr command is used to change file attributes on Linux systems that support extended attributes, such as ext2, ext3, ext4, btrfs, xfs, and others. File attributes are flags that modify the behavior of files and directories.

To prevent a file from being modified by all users in a Linux system, including the root account, the development team can use the chattr +i /app/conf/file command. This command will set the immutable attribute (+i) on the file /app/conf/file, which means that the file cannot be deleted, renamed, linked, appended, or written to by any user or process. To remove the immutable attribute, the development team can use the chattr -i /app/conf/file command. The statement C is correct.

The statements A, B, and D are incorrect because they do not prevent the file from being modified by all users. The chmod /app/conf/file command does not work because it requires an argument to specify the permissions to change. The setenforce /app/conf/file command does not work because it is used to change the SELinux mode, not file attributes. The chmod 0000 /app/conf/file command will remove all permissions from the file, but it can still be modified by the root account. References: [How to Use chattr Command in Linux]

**NEW QUESTION 24**

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

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

**Answer: D**

**Explanation:**

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

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

? [How to Use Ambassador Containers]

**NEW QUESTION 28**

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

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

**Answer: A**



#### Explanation:

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

#### NEW QUESTION 33

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

Output 1:

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

Output 2:

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

Output 3:

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

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

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

**Answer: D**

#### Explanation:

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

#### NEW QUESTION 37

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 41**

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 44**

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 48**

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 51**

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:**

This command will remove all containers, regardless of their state, by passing the IDs of all containers to the docker rm command. The docker ps -aq command will list the IDs of all containers, including the ones in an exited state, and the \$ ( ) syntax will substitute the output of the command as an argument for the docker rm command. This is a quick and easy way to clean up all containers, but it may also remove containers that are still needed or running.

#### References

? docker rm | Docker Docs - Docker Documentation, section "Remove all containers"

? Docker Remove Exited Containers | Easy methods. - Bobcares, section "For removing all exited containers"

#### NEW QUESTION 55

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. gitignore
- 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 60

A Linux administrator cloned an existing Linux server and built a new server from that clone. The administrator encountered the following error after booting the cloned server:

Device mismatch detected

The administrator performed the commands listed below to further troubleshoot and mount the missing filesystem:

```
#ls -al /dev/disk/by-uuid/  
total 0  
drwxr-xr-x 2 root 220 Jul 08:59 .  
drwxr-xr-x 2 root 160 Jul 08:59 ..  
lrwxrwxrwx 1 root 26 Jul 11:10 2251a54-6c14-9187-df8629373 -> ../../sdb  
lrwxrwxrwx 1 root 26 Jul 11:10 4211c54-2a13-7291-bd8629373 -> ../../sdc  
lrwxrwxrwx 1 root 26 Jul 11:10 3451b54-6d10-3561-ad8629373 -> ../../sdd
```

Which of the following should administrator use to resolve the device mismatch issue and mount the disk?

- A. mount disk by device-id
- B. fsck -A
- C. mount disk by-label
- D. mount disk by-blkid

**Answer:** A

#### Explanation:

The administrator should use the command mount disk by device-id to resolve the device mismatch issue and mount the disk. The issue is caused by the cloned server having a different device name for the disk than the original server. The output of blkid shows that the disk has the device name /dev/sdb1 on the cloned server, but the output of cat /etc/fstab shows that the disk is expected to have the device name /dev/sda1. The command mount disk by device-id will mount the disk by using its unique identifier (UUID) instead of its device name. The UUID can be obtained from the output of blkid or lsblk -f. The command will mount the disk to the specified mount point (/data) and resolve the issue. The other options are incorrect because they either do not mount the disk (fsck -A), do not use the correct identifier (mount disk by-label or mount disk by-blkid), or do not exist (mount disk by-blkid). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 318-319.

#### NEW QUESTION 61

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

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

**Answer:** A

#### Explanation:

The command systemctl status systemd-resolved.service will show the information about the service systemd-resolved.service. The systemctl command is a tool for managing system services and units. The status option displays the current status of a unit, such as active, inactive, or failed. The output also shows the unit description, loaded configuration, process ID, memory usage, and recent log messages. This command will show if the service systemd-resolved.service is running without any errors. This is the

correct command to use to accomplish the task. The other options are incorrect because they either perform different actions (enable, mask, or show) or do not show the status of the service (systemctl show systemd-resolved.service only shows the properties of the service, not the status). References: CompTIA Linux+



(XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, page 427.

**NEW QUESTION 66**

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 68**

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 71**

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

- A. rsync user@10.10.10.80: /tmp accounts.pdf
- B. scp accounts.pdf user@10.10.10.80:/tmp
- C. cp user@10.10.10. 80: /tmp accounts.pdf
- D. ssh accounts.pdf user@10.10.10.80: /tmp

**Answer:** B


**Explanation:**

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

**NEW QUESTION 76****DRAG DROP**

A new drive was recently added to a Linux system. Using the environment and tokens provided, complete the following tasks:

- Create an appropriate device label.
- Format and create an ext4 file system on the new partition. The current working directory is /.



- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

To create an appropriate device label, format and create an ext4 file system on the new partition, you can use the following commands:

? To create a GPT (GUID Partition Table) label on the new drive /dev/sdc, you can use the parted command with the -s option (for script mode), the device name (/dev/sdc), the mklable command, and the label type (gpt). The command is:

parted -s /dev/sdc mklable gpt

? To create a primary partition of 10 GB on the new drive /dev/sdc, you can use the parted command with the -s option, the device name (/dev/sdc), the mkpart command, the partition type (primary), the file system type (ext4), and the start and end points of the partition (1 and 10G). The command is:

parted -s /dev/sdc mkpart primary ext4 1 10G

? To format and create an ext4 file system on the new partition /dev/sdc1, you can use the mkfs command with the file system type (ext4) and the device name (/dev/sdc1). The command is:

mkfs.ext4 /dev/sdc1

You can verify that the new partition and file system have been created by using the lsblk command, which will list all block devices and their properties.

**NEW QUESTION 79**

What is the main objective when using Application Control?

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

**Answer: D**

**Explanation:**

The main objective when using Application Control is to ensure the security and privacy of information. Application Control is a security practice that blocks or restricts unauthorized applications from executing in ways that put data at risk. The control functions vary based on the business purpose of the specific application, but the main objective is to help ensure the privacy and security of data used by and transmitted between applications<sup>1</sup>. Application Control can also prevent malware, untrusted, or unwanted applications from running on the network, reducing the risks and costs associated with data breaches<sup>1</sup>. Application Control can also improve the overall network stability and performance by eliminating unnecessary or harmful applications<sup>1</sup>.

Application Control is not mainly used to filter out specific content, although it can be combined with other technologies such as URL filtering or content filtering to achieve that goal. Application Control is not mainly used to assist the firewall blade with handling traffic, although it can be integrated with firewall policies to enforce granular access rules based on applications. Application Control is not mainly used to see what users are doing, although it can provide visibility and reporting on application usage and activity.

**NEW QUESTION 80**

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<sup>2</sup>. References: 1: Git - git-fetch Documentation 2: Git Fetch | Atlassian Git Tutorial

**NEW QUESTION 83**

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 84**

A Linux system is failing to boot with the following error:

```
error: no such partitions
Entering rescue mode...
grub rescue>
```

Which of the following actions will resolve this issue? (Choose two.)

- A. Execute grub-install --root-directory=/mnt and reboot.
- B. Execute grub-install /dev/sdX and reboot.
- C. Interrupt the boot process in the GRUB menu and add rescue to the kernel line.
- D. Fix the partition modifying /etc/default/grub and reboot.
- E. Interrupt the boot process in the GRUB menu and add single to the kernel line.
- F. Boot the system on a LiveCD/ISO.

**Answer:** BF

**Explanation:**

The administrator should do the following two actions to resolve the issue:

? Boot the system on a LiveCD/ISO. This is necessary to access the system and repair the boot loader. A LiveCD/ISO is a bootable media that contains a Linux distribution that can run without installation. The administrator can boot the system from the LiveCD/ISO and mount the root partition of the system to a temporary directory, such as /mnt.

? Execute grub-install /dev/sdX and reboot. This will reinstall the GRUB boot loader to the disk device, where sdX is the device name of the disk, such as sda or sdb. The GRUB boot loader is a program that runs when the system is powered on and allows the user to choose which operating system or kernel to boot. The issue is caused by a corrupted or missing GRUB boot loader, which prevents the system from booting. The command grub-install will restore the GRUB boot loader and fix the issue.

The other options are incorrect because they either do not fix the boot loader (interrupt the boot process in the GRUB menu or fix the partition modifying /etc/default/grub) or do not use the correct syntax (grub-install --root-directory=/mnt instead of grub-install /dev/sdX or rescue or single instead of recovery in the GRUB

menu). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing the Linux Boot Process, pages 265-266.

**NEW QUESTION 85**

A Linux administrator needs to connect securely to a remote server in order to install application software. Which of the following commands would allow this connection?

- A. scp "ABC-key.pem" root@10.0.0.1
- B. sftp rooteiO.0.0.1
- C. telnet 10.0.0.1 80
- D. ssh -i "ABC-key.pem" root@10.0.0.1
- E. sftp "ABC-key.pem" root@10.0.0.1

**Answer:** D

**Explanation:**

The command ssh -i "ABC-key.pem" root@10.0.0.1 would allow the administrator to connect securely to the remote server in order to install application software. The ssh command is a tool for establishing secure and encrypted connections between remote systems. The -i option specifies the identity file that contains the private key for key-based authentication. The "ABC-key.pem" is the name of the identity file that contains the private key. The root@10.0.0.1 is the username and the IP address of the remote server. The command ssh -i "ABC-key.pem" root@10.0.0.1 will connect to the remote server using the private key and allow the administrator to install application software. This is the correct command to use to connect securely to the remote server. The other options are incorrect because they either do not use key-based authentication (sftp root@10.0.0.1 or telnet 10.0.0.1 80) or do not use the correct syntax for the command (scp "ABC-key.pem" root@10.0.0.1 instead of scp -i "ABC-key.pem" root@10.0.0.1 or sftp "ABC-key.pem" root@10.0.0.1 instead of sftp -i "ABC-key.pem" root@10.0.0.1). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Implementing Basic Security, page 513.

**NEW QUESTION 89**

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 94**

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 99**

An administrator accidentally deleted the /boot/vmlinuz file and must resolve the issue before the server is rebooted. Which of the following commands should the administrator use to identify the correct version of this file?

- A. rpm -qa | grep kernel; uname -a
- B. yum -y update; shutdown -r now
- C. cat /etc/centos-release; rpm -Uvh --nodeps
- D. telinit 1; restorecon -Rv /boot

**Answer:** A

**Explanation:**

The command rpm -qa | grep kernel lists all the installed kernel packages, and the command uname -a displays the current kernel version. These commands can help the administrator identify the correct version of the /boot/vmlinuz file, which is the kernel image file. The other options are not relevant or helpful for this task.

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

**NEW QUESTION 104**

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 -lh2 .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 109**

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 111**

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 115**

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 116**

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 118**

A Linux administrator is troubleshooting SSH connection issues from one of the workstations.

When users attempt to log in from the workstation to a server with the IP address 104.21.75.76, they receive the following message:

```
ssh: connect to host 104.21.75.76 port 22: Connection refused
```

The administrator reviews the information below:

**Workstation output 1:**

```
eth0: <BROADCAST,MULTICAST, UP, LOWER_UP> mtu 1500 qdisc mq state UP group default
link/ether 00:15:5d:e9:e9:fb brd 5.189.153.255 scope global eth0
inet 5.189.153.89/24 brd 5.189.153.255 scope global eth0
```

**Workstation output 2:**

```
default via 5.189.153.1 dev eth0
5.189.153.0/24 dev eth0 proto kernel scope link src 5.189.153.89
```

**Server output 1:**

target	prot	opt	source	destination
REJECT	tcp	--	101.68.78.194	0.0.0.0/0 tcp dpt:22 ctstate NEW, UNTRACKED reject-with icmp-port-unreachable
REJECT	tcp	--	222.186.180.130	0.0.0.0/0 tcp dpt:22 ctstate NEW, UNTRACKED reject-with icmp-port-unreachable
REJECT	tcp	--	104.131.1.39	0.0.0.0/0 tcp dpt:22 ctstate NEW, UNTRACKED reject-with icmp-port-unreachable
REJECT	tcp	--	68.183.196.11	0.0.0.0/0 tcp dpt:22 ctstate NEW, UNTRACKED reject-with icmp-port-unreachable
REJECT	tcp	--	5.189.153.89	0.0.0.0/0 tcp dpt:22 ctstate NEW, UNTRACKED reject-with icmp-port-unreachable
REJECT	tcp	--	41.93.32.148	0.0.0.0/0 tcp dpt:22 ctstate NEW, UNTRACKED reject-with icmp-port-unreachable

**Server output 2:**

```
sshd.service - OpenSSH server daemon
Loaded: loaded (/usr/lib/systemd/system/sshd.service; disabled; vendor preset: enabled)
Active: active (running) since Thu 2021-08-26 18:50:19 CEST; 2 weeks 5 days ago
```

**Server output 3:**

```
eth0: <BROADCAST, MULTICAST, UP, LOWER_UP> mtu 1500 qdisc mq state UP group default
link/ether 52:52:00:2a:bb:98 brd 104.21.75.255 scope global eth0
inet 104.21.75.76/24 brd 104.21.75.255 scope global eth0
```

**Server output 4:**

```
default via 104.21.75.254 dev eth0
104.21.75.0/24 dev eth0 proto kernel scope link src 104.21.75.76
```

Which of the following is causing the connectivity issue?

- A. The workstation has the wrong IP settings.
- B. The sshd service is disabled.
- C. The server's firewall is preventing connections from being made.
- D. The server has an incorrect default gateway configuration.

**Answer: C**

**Explanation:**

The server's firewall is preventing connections from being made, which is causing the connectivity issue. The output of iptables -L -n shows that the firewall is blocking all incoming traffic on port 22, which is the default port for SSH. The output of ssh -v user@104.21.75.76 shows that the connection is refused by the server. To resolve the issue, the administrator needs to allow port 22 on the firewall. The other options are incorrect because they are not supported by the

outputs. The workstation has the correct IP settings, as shown by the output of `ip addr show`. The `sshd` service is enabled and running, as shown by the output of `systemctl status sshd`. The server has the correct default gateway configuration, as shown by the output of `ip route show`. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, pages 406-407.

#### NEW QUESTION 121

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

Output 1:

\* `httpd.service` = The Apache HTTPD Server

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

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

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

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

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

**Answer:** CD

#### Explanation:

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

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

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

#### NEW QUESTION 124

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

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

**Answer:** B

#### Explanation:

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

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

#### NEW QUESTION 126

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

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

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 131

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

```
# grep joe /etc/passwd /etc/shadow
/etc/passwd:joe:x:1001:1001::/home/joe:/bin/nologin
/etc/shadow:joe:$6$3uOw6qWx9876jGhgKJsdFh987634534voj.:18883:0:99999:7:::
```

Which of the following command would resolve the issue?



- A. usermod -s /bin/bash joe
- B. pam\_tally2 -u joe -r
- C. passwd -u joe
- D. chage -E 90 joe

**Answer: B**

**Explanation:**

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

**NEW QUESTION 134**

The security team has identified a web service that is running with elevated privileges. A Linux administrator is working to change the `systemd` service file to meet security compliance standards. Given the following output:

```
[Unit]
Description=CompTIA server daemon
Documentation=man:webserver(8) man:webserver_config(5)
After=network.target

[Service]
Type=notify
EnvironmentFile=/etc/webserver/config
ExecStart=/usr/sbin/webserver -D $OPTIONS
ExecReload=/bin/kill -HUP $MAINPID
KillMode=process
Restart=on-failure
RestartSec=42s

[Install]
WantedBy=multi-user.target
```

Which of the following remediation steps will prevent the web service from running as a privileged user?

- A. Removing the `ExecStarWusr/sbin/webserver -D SOPTIONS` from the service file
- B. Updating the Environment File line in the `[Service]` section to `/home/websevice/config`
- C. Adding the `User=websevice` to the `[Service]` section of the service file
- D. Changing the `multi-user.target` in the `[Install]` section to `basic.target`

**Answer: C**

**Explanation:**

The remediation step that will prevent the web service from running as a privileged user is adding the `User=websevice` to the `[Service]` section of the service file. The service file is a configuration file that defines the properties and behavior of a `systemd` service. The `systemd` is a system and service manager that controls the startup and operation of Linux systems. The service file contains various sections and options that specify how the service should be started, stopped, and managed. The `[Service]` section defines how the service should be executed and what commands should be run. The `User` option specifies the user name or ID that the service should run as. The `websevice` is the name of the user that the administrator wants to run the web service as. The administrator should add the `User=websevice` to the `[Service]` section of the service file, which will prevent the web service from running as a privileged user, such as `root`, and improve the security of the system. This is the correct remediation step to use to prevent the web service from running as a privileged user. The other options are incorrect because they either do not change the user that the service runs as (removing the `ExecStart=/usr/sbin/webserver -D OPTIONS` from the service file or updating the `EnvironmentFile` line in the `[Service]` section to `/home/websevice/config`) or do not affect the user that the service runs as (changing the `multi-user.target` in the `[Install]` section to `basic.target`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing System Services, page 458.

**NEW QUESTION 137**

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

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

**Answer: C**

**Explanation:**

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

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

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

The other commands are incorrect because:

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

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



? sed '/UID/' /etc/passwd < 500 will use sed to edit the /etc/passwd file and replace any line that contains UID with 500, not list the users with UID greater than 500. The < 500 part will redirect the input from a file named 500, not compare with the UID.

References:

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

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

#### NEW QUESTION 139

An application developer received a file with the following content:

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

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

```
RUN make /app
```

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

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

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

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

**Answer:** A

#### Explanation:

The docker build command is used to build an image from a Dockerfile and a context<sup>1</sup>. The Dockerfile is a text file that contains the instructions for creating the image, and the context is a set of files that can be used in the image creation process<sup>1</sup>. The file that the developer received is an example of a Dockerfile.

The -t option is used to specify a name and an optional tag for the image<sup>1</sup>. The name and tag are separated by a colon (:), and the tag is usually used to indicate the version of the image<sup>2</sup>. For example, -t myimage:1.0 means that the image will be named myimage and tagged as 1.0.

The last argument of the docker build command is the path to the context, which can be a local directory or a URL<sup>1</sup>. The dot (.) means that the current working directory is the context<sup>2</sup>. Therefore, docker build -t myimage:1.0 . means that the image will be built from the Dockerfile and the files in the current working directory, and it will be named myimage and tagged as 1.0.

#### NEW QUESTION 141

A cloud engineer needs to remove all dangling images and delete all the images that do not have an associated container. Which of the following commands will help to accomplish this task?

- A. docker images prune -a
- B. docker push images -a
- C. docker rmi -a images
- D. docker images rmi --all

**Answer:** A

#### Explanation:

The command docker images prune -a will help to remove all dangling images and delete all the images that do not have an associated container.

The docker command is a tool for managing Docker containers and images.

The images subcommand operates on images. The prune option removes unused images.

The -a option removes all images, not just dangling ones. A dangling image is an image that is not tagged and is not referenced by any container. This command will accomplish the task of cleaning up the unused images. The other options are incorrect because they either do not exist (docker push images -a or docker images rmi --all) or do not remove images (docker rmi -a images only removes images that match the name or ID of “images”). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 567.

#### NEW QUESTION 142

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

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

**Answer:** C

#### Explanation:

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

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

#### NEW QUESTION 146

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

- A. snap list

- B. snap find
- C. snap install
- D. snap try

**Answer:** A

**Explanation:**

The snap list command is used to display the installed snaps on the system<sup>1</sup>. Snaps are self-contained software packages that can be installed and updated across different Linux distributions<sup>2</sup>. The snap list command shows the name, version, revision, developer and notes of each snap<sup>1</sup>. The snap find command is used to search for snaps in the Snap Store, which is an online repository of snaps<sup>2</sup>. The snap install command is used to install snaps from the Snap Store or from a local file<sup>2</sup>. The snap try command is used to test a snap without installing it, by mounting a directory that contains the snap files<sup>2</sup>. These commands are not useful for verifying if a package was installed using a snap.

**NEW QUESTION 148**

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 152**

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

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

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

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

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

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

**Answer:** B

**Explanation:**

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

? CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Networking, Section: Troubleshooting Network Issues, Page 359.

? Linux+ (Plus) Certification, Exam Objectives: 4.3 Given a scenario, troubleshoot and resolve basic network configuration and connectivity issues.

**NEW QUESTION 153**

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

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

Which of the following commands would resolve the security issue?

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

**Answer:** D

**Explanation:**

This command will use the chsh tool to change the login shell of the users www-data and db to /bin/false, which means they will not be able to log in to the system<sup>1</sup>. This will prevent unauthorized access attempts and improve security.

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

**NEW QUESTION 157**

As part of the requirements for installing a new application, the swappiness parameter needs to be changed to O. This change needs to persist across re-boots and be applied immediately. A Linux systems administrator is performing this change. Which of the following steps should the administrator complete to accomplish this task?

- A. echo "v
- B. swappiness—()" >> /etc/sysctl . conf && sysctl —p
- C. echo "vr
- D. >> / proc/meminfo && sysctl —a
- E. sysctl —v >> / proc/meminfo & & echo "v
- F. swapiness=0"
- G. sysctl —h "v
- H. swapiness—O" && echo / etc/vmswapiness

**Answer:** A

**Explanation:**

To change the swappiness parameter to 0 and make it persistent across reboots and applied immediately, the administrator can perform the following steps:

? Append the line vm.swappiness=0 to the file /etc/sysctl.conf using echo

"vm.swappiness=0" >> /etc/sysctl.conf (A). This will set the swappiness parameter to 0 for future boots.

? Reload the sysctl configuration using sysctl -p (A). This will apply the changes to the current system without rebooting. The other commands will not achieve this task, but either write to a wrong file, use a wrong option, or have a syntax error. References:

? [CompTIA Linux+ Study Guide], Chapter 8: Optimizing Linux Performance, Section: Tuning Kernel Parameters with sysctl

? [How to Change Swappiness in Linux]

**NEW QUESTION 159**

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

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

**Answer:** AB

**Explanation:**

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

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

the user from the Kerberos key distribution center (KDC). The user needs to enter their password or use a keytab file to authenticate<sup>1</sup>.

? klist: This command lists the cached tickets, including the TGT and any service tickets, for the user. It also shows the expiration time and flags for each ticket<sup>2</sup>.

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

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

\$ klist

Ticket cache: FILE:/tmp/krb5cc\_1000 Default principal: username@REALM

Valid starting Expires Service principal

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

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

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

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

**NEW QUESTION 162**

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 164**

A systems administrator created a new Docker image called test. After building the image, the administrator forgot to version the release. Which of the following will allow the administrator to assign the v1 version to the image?

- A. docker image save test test:v1
- B. docker image build test:vl
- C. docker image tag test test:vl
- D. docker image version test:v1

**Answer:** C

**Explanation:**

The docker image tag test test:v1 command can be used to assign the v1 version to the image called test. This command creates a new tag for the existing

image, without changing the original image. The docker image save test:v1 command would save the image to a file, not assign a version. The docker image build test:vl command is invalid, as vl is not a valid version number. The docker image version test:v1 command does not exist. References: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 16: Virtualization and Cloud Technologies, page 500.

**NEW QUESTION 166**

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

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

**Answer:** A

**Explanation:**

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

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

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

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

**NEW QUESTION 170**

A cloud engineer needs to launch a container named web-01 in background mode. Which of the following commands will accomplish this task?"

- A. docker builder -f —name web-01 httpd
- B. docker load --name web-01 httpd
- C. docker ps -a --name web-01 httpd
- D. docker run -d --name web-01 httpd

**Answer:** D

**Explanation:**

The docker run -d --name web-01 httpd command will launch a container named web-01 in background mode. This command will create and start a new container from the httpd image, assign it the name web-01, and run it in detached mode (-d), which means the container will run in the background without attaching to the current terminal. The docker builder -f --name web-01 httpd command is invalid, as builder is not a valid docker command, and -f and --name are not valid options for docker build. The docker load --name web-01 httpd command is invalid, as load does not accept a --name option, and httpd is not a valid file name for load. The docker ps -a --name web-01 httpd command is invalid, as ps does not accept a --name option, and httpd is not a valid filter for ps. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Virtualization and Cloud Technologies, page 499.

**NEW QUESTION 175**

A Linux administrator generated a list of users who have root-level command-line access to the Linux server to meet an audit requirement. The administrator analyzes the following /etc/passwd and /etc/sudoers files:

```
$ cat /etc/passwd
```

```
root:x: 0:0: /home/root: /bin/bash lee: x: 500: 500: /home/lee:/bin/tcsh
```

```
mallory:x: 501:501: /root:/bin/bash
```

```
eve:x: 502: 502: /home/eve:/bin/nologin carl:x:0:503: /home/carl:/bin/sh
```

```
bob:x: 504: 504: : /home/bob:/bin/ksh
```

```
alice:x: 505:505: /home/alice:/bin/rsh
```

```
$ cat /etc/sudoers
```

```
Cmnd_Alias SHELLS = /bin/tcsh, /bin/sh, /bin/bash Cmnd_Alias SYSADMIN = /usr/sbin/tcpdump
```

```
ALL = (ALL) ALL
```

```
ALL = NOPASSWD: SYSADMIN
```

Which of the following users, in addition to the root user, should be listed in the audit report as having root-level command-line access? (Select two).

- A. Carl
- B. Lee
- C. Mallory
- D. Eve
- E. Bob
- F. Alice

**Answer:** AC

**Explanation:**

The users who have root-level command-line access are those who have either the same user ID (UID) as root, which is 0, or the ability to run commands as root using sudo. Based on the /etc/passwd and /etc/sudoers files, the users who meet these criteria are:

? Carl: Carl has the same UID as root, which is 0, as shown in the /etc/passwd file.

This means that Carl can log in as root and execute any command with root privileges1

? Mallory: Mallory has the ability to run commands as root using sudo, as shown in the /etc/sudoers file. The line ALL = (ALL) ALL means that any user can run any command as any other user, including root, by using sudo. Mallory can also use the root shell /bin/bash as her login shell, as shown in the /etc/passwd file2

Therefore, the correct answer is A and C. Lee, Eve, Bob, and Alice do not have root-level command-line access because they have different UIDs from root and they cannot use sudo to run commands as root. Lee can only use sudo to run the commands listed in the Cmnd\_Alias SHELLS, which are /bin/tcsh, /bin/sh, and /bin/bash. Eve cannot log in at all because her login shell is /bin/nologin. Bob and Alice can only use sudo to run the command /usr/sbin/tcpdump without a password, as specified by the Cmnd\_Alias SYSADMIN and the line ALL = NOPASSWD: SYSADMIN2

**NEW QUESTION 180**



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

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

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

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

Which of the following commands will permanently resolve the issue?

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

**Answer: D**

**Explanation:**

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

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

**NEW QUESTION 185**

Based on an organization's new cybersecurity policies, an administrator has been instructed to ensure that, by default, all new users and groups that are created fall within the specified values below.

```
# Min/max values for automatic uid selection in useradd
#
UID_MIN 1000
UID_MAX 60000
# Min/max values for automatic gid selection in groupadd
#
GID_MIN 1000
GID_MAX 60000
```

To which of the following configuration files will the required changes need to be made?

- A. /etc/login.defs
- B. /etc/security/limits.conf
- C. /etc/default/useradd
- D. /etc/profile

**Answer: A**

**Explanation:**

The required changes need to be made to the /etc/login.defs configuration file. The /etc/login.defs file defines the default values for user and group IDs, passwords, shells, and other parameters for user and group creation. The file contains the directives UID\_MIN, UID\_MAX, GID\_MIN, and GID\_MAX, which set the minimum and maximum values for automatic user and group ID selection. The administrator can edit this file and change the values to match the organization's new cybersecurity policies. This is the correct file to modify to accomplish the task. The other options are incorrect because they either do not affect the user and group IDs (/etc/security/limits.conf or /etc/profile) or do not set the default values (/etc/default/useradd). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Users and Groups, page 463.

**NEW QUESTION 189**

.....

## Thank You for Trying Our Product

### We offer two products:

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

2nd - Questions and Answers in PDF Format

### XK0-005 Practice Exam Features:

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

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