PassQueen
a.k.a dumpsqueen.com

CompTIA Linux+ Exam

CompTIA XK0-005

Version Demo

Total Demo Questions: 52

Total Premium Questions: 520

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, System Management	239
Topic 2, Security	92
Topic 3, Scripting, Containers, and Automation	86
Topic 4, Troubleshooting	103
Total	520

QUESTION NO: 1

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

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: B F

Explanation:

Booting the system on a LiveCD/ISO provides an independent Linux environment when the installed operating system cannot start. From that environment, an administrator can identify and mount the installed system's root and boot-related filesystems, then perform boot-loader recovery work without relying on the damaged installation. This is the appropriate first recovery action for a system whose GRUB boot loader is missing or corrupted.

Executing `grub-install /dev/sdX` and rebooting reinstalls GRUB to the boot disk, where `/dev/sdX` must be replaced with the actual target disk, such as `/dev/sda`. In normal recovery practice, this command is run with the installed system mounted and used as the recovery environment, commonly through a chroot; the installation's GRUB configuration should then be regenerated when needed. Reinstalling the boot loader restores the code and files needed to transfer firmware control to GRUB and proceed with kernel booting. GNU documents `grub-install` as the utility used to install GRUB on a device, and its manual describes the required installation process: GNU GRUB: Invoking grub-install. A practical mounted-system recovery workflow is also documented by ArchWiki: chroot.

QUESTION NO: 2

A systems administrator is trying to track down a rogue process that has a TCP listener on a network interface for remote command-and-control instructions.

Which of the following commands should the systems administrator use to generate a list of rogue process names? (Select two).

A. netstat -antp | grep LISTEN

B. lsof -iTCP | grep LISTEN

C. lsof -i:22 | grep TCP

D. netstat -a | grep TCP

E. nmap -p1-65535 | grep -i tcp

F. nmap -sS 0.0.0.0/0

ANSWER: A B

Explanation:

`netstat -antp | grep LISTEN` is appropriate because the TCP-related flags limit the displayed sockets to TCP, while `-a` includes listening sockets and `-p` requests the PID and program name associated with each socket. Filtering for the `LISTEN` state focuses the result on processes accepting inbound TCP connections, which is the relevant condition when

investigating a potential command-and-control listener. On many systems, the administrator must run this command with sufficient privileges to see program ownership information for processes owned by other users.

`lsof -iTCP | grep LISTEN` also directly supports the investigation. `lsof` associates open network files and sockets with the owning command and process identifier; the `-iTCP` selection restricts output to TCP network connections. Its output includes the TCP state, so filtering for `LISTEN` produces a concise list of processes that currently have TCP listening sockets. The command name and PID can then be correlated with service configuration, package ownership, process ancestry, and executable paths to identify suspicious listeners. See the netstat manual page and the lsof manual page.

QUESTION NO: 3

An administrator is troubleshooting a database service outage that was reported by a monitoring system. Given the following output:

$ systemctl status mariadb

Oct 20 16:40:45 comptia systemd[1]: mariadb.service: Main process exited, code=killed, status=9/KILL

Oct 20 16:40:45 comptia systemd[1]: mariadb.service: Failed with result ' signal ' .

Oct 20 16:40:50 comptia systemd[1]: Stopped MariaDB 10.3 database server.

$ dmesg

[1061.491433] oom-kill:constraint=CONSTRAINT_NONE,nodemask=(null),cpuset=/,mems_allowed=0,global_oom, task_memcg=/system.slice/mariadb.service, task=mysqld,pid=1981,uid=27

[1061.491453] Out of memory: Killed process 1981 (mysqld) total-vm:330668kB, anon-rss:31316kB, file-rss:OkB, shmem-rss:OkB, UID:27 pgtables:324kB oom_score_adj:0

Which of the following is the reason for the outage?

A. The administrator sent a kill signal to the database

B. The server is missing the DMA bus

C. The database cannot write anything else to the storage

D. The server does not have enough physical memory

ANSWER: D

Explanation:

The server does not have enough physical memory. The kernel log explicitly records a global out-of-memory event with `global_oom`, followed by `Out of memory: Killed process 1981 (mysqld)`. When Linux cannot satisfy memory-allocation requests because available RAM (and any usable swap) is exhausted, its out-of-memory killer selects a process to terminate and sends it SIGKILL. In this case, the selected process was `mysqld`, which is the MariaDB database-server process. This directly corresponds to the systemd status showing that the MariaDB service's main process exited with status `9/KILL`; signal 9 is SIGKILL. Therefore, the service outage resulted from memory exhaustion causing the kernel to terminate MariaDB, rather than from a normal administrative service stop. Administrators should investigate overall memory pressure, concurrent workloads, MariaDB memory configuration, available swap, and any applicable systemd memory limits before restarting the service. The Linux kernel documentation describes the out-of-memory killer and its process-selection behavior in the Linux kernel OOM documentation. Red Hat also documents identifying OOM-killer events through kernel logs in its OOM troubleshooting guidance.

QUESTION NO: 4

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` is correct because `git clone` creates a new local copy of an existing remote Git repository. It contacts the repository at the supplied HTTPS URL, downloads its commits, branches, tags, and tracked files, and initializes a local working directory. By default, Git derives the destination directory name from the repository name, so this command ordinarily creates a directory named `project` in the current directory. It also configures a remote named `origin`, allowing later synchronization commands such as `git fetch`, `git pull`, and `git push` to use that saved remote location. HTTPS cloning may prompt the engineer for credentials or a personal access token if the company repository is private. This is the standard workflow for obtaining a repository that does not already exist on the local system. Git's official documentation describes `git clone` as cloning a repository into a new directory and setting up remote-tracking branches and a remote configuration. See the Git clone documentation and the Pro Git guide to getting a Git repository.

QUESTION NO: 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: A E

Explanation:

For a systemd mount unit, the unit filename must correspond exactly to the mount path specified by `Where=`, after systemd path escaping has been applied. The desired path is `/home/user1/My Documents`, so its corresponding mount-unit name is `home-user1-My\x20Documents.mount`. The leading slash is omitted, path separators become hyphens, and the space is represented as the hexadecimal escape sequence `\x20`. Therefore, renaming the unit to `home-user1-My\x20Documents.mount` establishes the required filename-to-mountpoint relationship.

The mount path itself also needs valid systemd escaping in the unit's `Where=` directive. Using `Where=/home/user1/My\x20Documents` preserves the intended directory name, including its uppercase letters and embedded space, while allowing systemd to parse the value correctly. After applying these corrections, the administrator should reload unit definitions with `systemctl daemon-reload` before enabling or starting the mount unit. The systemd mount-unit documentation explicitly requires the unit name to be derived from the escaped `Where=` path, and the systemd escape utility documents this path-escaping convention. See systemd.mount and systemd-escape.

QUESTION NO: 6

A systems administrator engineer wants to configure a Linux server's time zone to America/Chicago. The engineer verifies that the server's time zone is configured for Asia/Tokyo:

Which of the following commands should the engineer execute to make this change?

A. date configure -set-timezone America/Chicago

B. set-timezone -config America/Chicago

C. time date set America/Chicago

D. timedatectl set-timezone America/Chicago

ANSWER: D

Explanation:

`timedatectl set-timezone America/Chicago` is the appropriate command on a systemd-based Linux distribution. The `timedatectl` utility manages system clock and time-zone settings through the systemd timedated service. Its `set-timezone` subcommand accepts an IANA time-zone identifier, such as `America/Chicago`, and updates the system-wide local time-zone configuration. America/Chicago is a valid zone name from the system time-zone database and automatically accounts for the applicable Central Time offset and daylight-saving-time transitions. The command does not manually alter the current clock value; instead, it changes how the system interprets and displays local time relative to the underlying system time. Administrators can confirm the resulting configuration with `timedatectl status` or simply `timedatectl`, which reports the configured local time zone. This approach is preferred on modern Linux systems using systemd because it uses the supported service interface and maintains the standard system time-zone configuration. See the timedatectl documentation and the IANA Time Zone Database.

QUESTION NO: 7

An administrator needs to make an application change via a script that must be run only in console mode. Which of the following best represents the sequence the administrator should execute to accomplish this task?

A. systemctl isolate multi-user.target; sh script.sh; systemctl isolate graphical.target

B. systemctl isolate graphical.target; sh script.sh; systemctl isolate multi-user.target

C. sh script.sh; systemctl isolate multi-user.target; systemctl isolate graphical.target

D. systemctl isolate multi-user.target; systemctl isolate graphical.target; sh script.sh

ANSWER: A

Explanation:

The sequence `systemctl isolate multi-user.target; sh script.sh; systemctl isolate graphical.target` is correct because it first transitions the system into the systemd `multi-user.target`, which provides the standard non-graphical, multiuser console environment. The administrator can then execute `script.sh` while the system is operating in that console-oriented target. After the application change completes, isolating `graphical.target` restores the normal graphical environment.

`systemctl isolate` changes the currently active target immediately: systemd starts the requested target and its required dependencies while stopping units that are not part of the new target, subject to unit dependency and isolation rules. This makes it appropriate when the environment must be changed temporarily without modifying the configured default boot target or rebooting the host. The semicolon separators represent running the commands sequentially; in an operational script, using appropriate error handling would ensure the graphical target is restored only after the maintenance action has completed as intended.

See the systemctl documentation for the behavior of `isolate`, and the systemd special targets documentation for descriptions of `multi-user.target` and `graphical.target`.

QUESTION NO: 8

Which of the following actions are considered good security practices when hardening a Linux server? (Select two).

A. Renaming the root account to something else

B. Removing unnecessary packages

C. Changing the default shell to /bin/csh

D. Disabling public key authentication

E. Disabling the SSH root login possibility

F. Changing the permissions on the root filesystem to 600

ANSWER: B E

Explanation:

Removing unnecessary packages is a core Linux-server hardening practice because every installed service, library, utility, and package can potentially introduce vulnerabilities, insecure defaults, or configuration work that must be maintained. Applying a minimal-installation approach reduces the system's attack surface and helps administrators focus patching, monitoring, and access control on software that is genuinely required for the server's role. The Linux+ exam objectives include applying hardening techniques such as reducing the attack surface and managing services securely.

Disabling the SSH root login possibility is also a sound practice. SSH access should be attributable to an individual administrative account, with privilege elevation performed only when necessary through controlled mechanisms such as `sudo`. Setting `PermitRootLogin no` in the OpenSSH server configuration prevents direct remote authentication as root, reducing exposure to attacks against the universally known root account and improving accountability in logs. This control should be paired with strong authentication and appropriate administrative-access policies. See the CompTIA exam objectives and the OpenSSH sshd_config documentation for PermitRootLogin.

QUESTION NO: 9 - (SIMULATION)

As a Systems Administrator, to reduce disk space, you were tasked to create a shell script that does the following:

Add relevant content to /tmp/script.sh, so that it finds and compresses rotated files in /var/log without recursion.

INSTRUCTIONS

Fill the blanks to build a script that performs the actual compression of rotated log files.

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

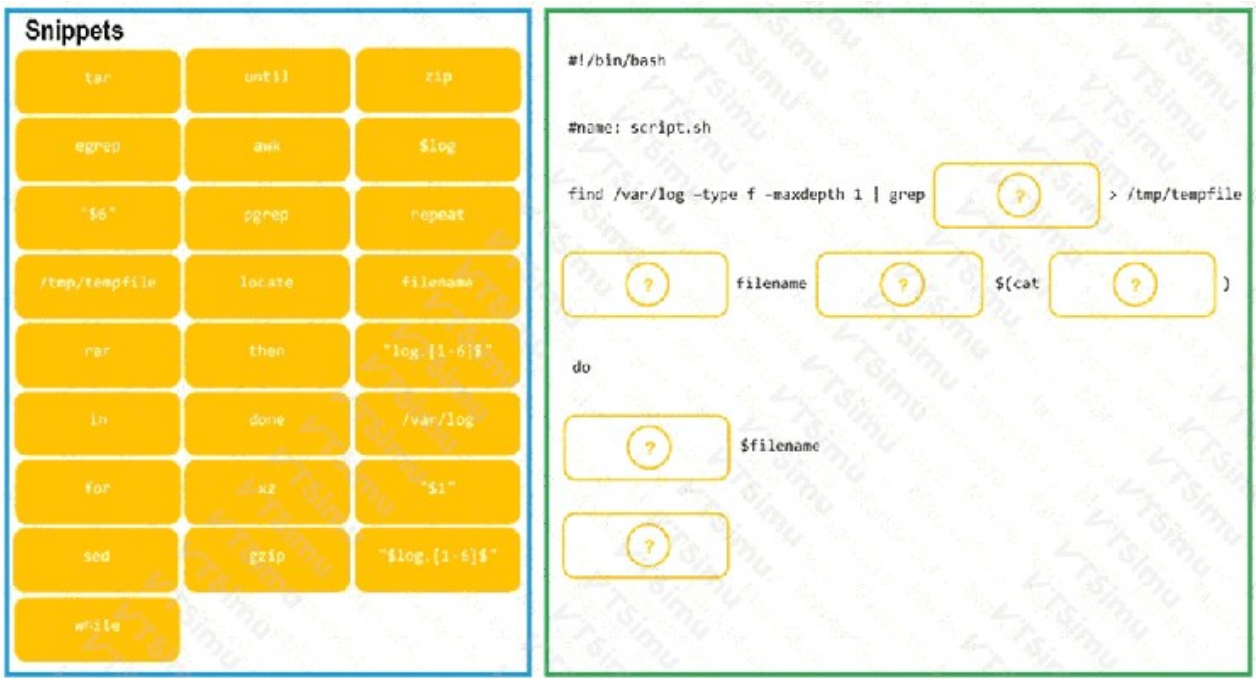

ANSWER: See the explanation for the answer

Explanation:

Fill the blanks so that the script is:

Use these snippets in order:

`-maxdepth 1` restricts `find` to files directly in `/var/log`, and the loop runs `gzip` against each matching rotated log filename.

QUESTION NO: 10

A developer reported an incident involving the application configuration file /etc/httpd/conf/httpd.conf that is missing from the server. Which of the following identifies the RPM package that installed the configuration file?

A. rpm -qf /etc/httpd/conf/httpd.conf

B. rpm -ql /etc/httpd/conf/httpd.conf

C. rpm —query /etc/httpd/conf/httpd.conf

D. rpm -q /etc/httpd/conf/httpd.conf

ANSWER: A

Explanation:

`rpm -qf /etc/httpd/conf/httpd.conf` is the appropriate command because the `-f` (or `--file`) query option asks RPM which installed package owns a specified filesystem path. RPM answers this by consulting its local RPM database, which records the files installed by each package. Consequently, this query can identify the owning package even when the referenced configuration file is currently missing from the filesystem, provided the RPM database entry remains installed and intact. This is particularly useful during incident response: an administrator can establish the package responsible for the expected file, then use the package name to inspect package metadata or reinstall the package if recovery is required. The command must be given the complete path to the expected file, as shown, so RPM can match that path against its recorded

package file list. The RPM manual documents `--file` as querying the package that owns a file; see the rpm(8) manual page. For additional RPM package-management context, see the Fedora DNF documentation.

QUESTION NO: 11

A Linux administrator needs to make the `developers` branch current with changes from the remote repository while preserving a linear local commit history. Which of the following commands can be used to accomplish this task? (Choose two.)

A. git fetch origin

B. git rebase origin/developers

C. git reset --hard origin/developers

D. git tag developers

E. git init origin/developers

ANSWER: A B

Explanation:

`git fetch origin` downloads updated references and commits from the remote named `origin` without immediately modifying the currently checked-out branch. The administrator can then review the fetched changes before integrating them.

`git rebase origin/developers` reapplies local commits from the current branch on top of the updated remote-tracking branch. When the local branch is `developers`, this incorporates the remote changes while avoiding a merge commit, resulting in a linear history. Any conflicts encountered during the rebase must be resolved before continuing with `git rebase --continue`. Git documents remote retrieval in the git-fetch reference and commit rebasing in the git-rebase reference.

QUESTION NO: 12

A Linux systems administrator wants to validate the network interfaces configuration and status on a Linux server. Which of the following commands should the administrator use to display the information for interface eth3 including status, IP address, subnet mask, and configuration?

A. ip list interface eth3

B. ip addr show dev eth3

C. ip show eth3

D. ip config eth3

ANSWER: B

Explanation:

The command `ip addr show dev eth3` is the appropriate modern Linux command for examining address-related configuration for the specific network interface named `eth3`. It displays the interface name and flags, operational state, link-layer (MAC) address, and the IPv4 and IPv6 addresses assigned to that interface. IPv4 addresses are shown in CIDR notation, such as `192.0.2.10/24`; the prefix length (`/24`) provides the subnet-mask information needed to determine the interface's local network. This command is especially useful for validating whether an interface has received or been assigned the intended addresses and whether it is administratively enabled and operationally active. The `ip` utility is part of the iproute2 suite and is the standard replacement for legacy interface-inspection workflows that used `ifconfig`. To focus the output on one device, the `dev eth3` selector is used with the `ip address show` object and action. For syntax and output details, consult the ip-address(8) manual page and the ip(8) manual page.

QUESTION NO: 13

The users of a Linux system are unable to use one of the application filesystems. The following outputs have been provided:

bash

$ cd /app

$ touch file

touch: cannot touch 'file': Readonly file system

Output 2

/dev/sdcl on /app type ext4 (ro,relatime,seclabel,data=ordered)

Output 3

/dev/sdcl /app ext4 defaults 0 0

Output 4

[302.048075] Buffer I/O error on dev sdcl, logical block 0, async page read

[302.048490] EXT4-fs (sdcl): Attempt to read block from filesystem resulted in short read while trying to re-open /dev/sdcl

Which of the following actions will resolve this issue?

A. umount /app fsck -y /dev/sdcl mount /app

B. xfs_repair /dev/sdcl mount -o rw,remount /app

C. umount /app xfs_repair /dev/sdcl mount /app

D. fsck -y /dev/sdcl mount -o rw,remount /app

ANSWER: A

Explanation:

`umount /app fsck -y /dev/sdcl mount /app` is the appropriate recovery sequence for the ext4 filesystem. The mount output confirms that `/app` uses ext4 and is currently mounted read-only. The kernel messages show I/O failures and an ext4 short read, conditions that can leave the filesystem unavailable for writes and may require filesystem consistency repairs. First, unmounting `/app` ensures that no filesystem metadata is being actively changed while it is checked. The `fsck -y /dev/sdcl` command then invokes the suitable filesystem checker for the ext4 device and automatically accepts repairs needed to restore consistency. After a successful check, `mount /app` mounts the filesystem again using its configured `/etc/fstab` entry, whose `defaults` option permits normal read-write mounting. The fsck documentation specifies that filesystems should generally be unmounted before checking or repair, because checking a mounted filesystem can produce unreliable results or cause damage. See the fsck manual page and the Linux kernel ext4 documentation. Because the reported errors are I/O-related, the storage device and its connections should also be investigated after service is restored.

QUESTION NO: 14

A systems administrator is notified that the mysqld process stopped unexpectedly. The systems administrator issues the following command:

sudo grep –i -r ‘out of memory’ /var/log

The output of the command shows the following:

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

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

A. free -h

B. nc -v 127.0.0.1 3306

C. renice -15 $(pidof mysql)

D. lsblk

E. killall -15

F. vmstat -a 1 4

ANSWER: A F

Explanation:

The kernel message confirms that the out-of-memory killer selected and terminated the database process because the system experienced memory exhaustion. The appropriate next troubleshooting steps are to establish the current memory and swap capacity and to observe memory-related virtual-memory activity over a short interval. `free -h` presents RAM and swap totals, current usage, cache, and available memory in human-readable units. This helps determine whether memory or swap remains constrained and provides an immediate baseline after the event. The `vmstat -a 1 4` command collects four one-second samples of system activity, including runnable and blocked processes, swap activity, I/O, CPU utilization, and memory information. In particular, recurring swap-in or swap-out activity and sustained resource pressure can support a diagnosis of insufficient memory or workload contention. Taking several samples is more useful than relying only on a static snapshot because it reveals whether pressure is ongoing. These commands gather the evidence needed to determine whether to investigate memory-consuming workloads, database configuration, available RAM, or swap configuration before making operational changes. See the free manual page and the vmstat manual page.

QUESTION NO: 15

A Linux administrator needs to check the content of a log file that is appending data as the file grows. Which of the following commands should the administrator use to accomplish this task?

A. tail -f log.txt

B. tail -v log.txt

C. tail -c log.txt

D. tail log.txt

ANSWER: A

Explanation:

The correct command is `tail -f log.txt`. The `tail` utility normally displays the end of a file, which is useful for reviewing the most recent log entries. Its `-f` (follow) option keeps the command running after the initial output and continuously displays additional data as it is appended to the file. This enables an administrator to observe events in near real time without repeatedly rerunning a command, making it particularly useful while troubleshooting a service, monitoring application activity, or watching system log messages.

For example, running `tail -f /var/log/syslog` or an application-specific log lets the administrator see new lines as the relevant daemon or application writes them. The process continues until it is interrupted, typically with `Ctrl+C`. GNU Coreutils documents that `--follow`, also available as `-f`, outputs appended data as the file grows. For logs that may be rotated and recreated, administrators can use `tail -F`, which follows the filename across replacement files; however, `tail -f log.txt` directly satisfies the stated requirement to follow data appended to the growing file. See the GNU Coreutils tail documentation and the tail(1) Linux manual page.

QUESTION NO: 16

A Linux administrator is creating a new sudo profile for the accounting user. Which of the following should be added by the administrator to the sudo

configuration file so that the accounting user can run / opt/ acc/ report as root?

A. accounting localhost=/opt/acc/report

B. accounting ALL=/opt/acc/report

C. %accounting ALL=(ALL) NOPASSWD: /opt/acc/report

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

ANSWER: B

Explanation:

`accounting ALL=/opt/acc/report` is the appropriate sudoers entry for the stated requirement. In sudoers syntax, the first field identifies the user, so `accounting` applies to the user named accounting. The `ALL` host field permits this rule to apply regardless of the host on which the sudoers file is used. The command specification then limits sudo access to the absolute path `/opt/acc/report`, following the principle of least privilege.

Most importantly, when a sudoers rule does not explicitly include a Runas specification such as `(user)`, sudo uses its default Runas user. That default is `root`. Therefore, invoking `sudo /opt/acc/report` under this rule runs the report program as root while retaining the normal sudo authentication behavior. No `NOPASSWD` tag is needed because the question requires root execution, not passwordless execution. The sudoers manual documents that the default Runas user is root and describes the user, host, runas, and command components of a rule. Administrators should edit this configuration with `visudo` to validate syntax before saving. See the sudoers manual and the visudo manual.

QUESTION NO: 17 - (DRAG DROP)

As a Systems Administrator, to reduce disk space, you were tasked to create a shell script that does the following:

Add relevant content to /tmp/script.sh, so that it finds and compresses rotated files in /var/log without recursion.

INSTRUCTIONS

Fill the blanks to build a script that performs the actual compression of rotated log files.

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

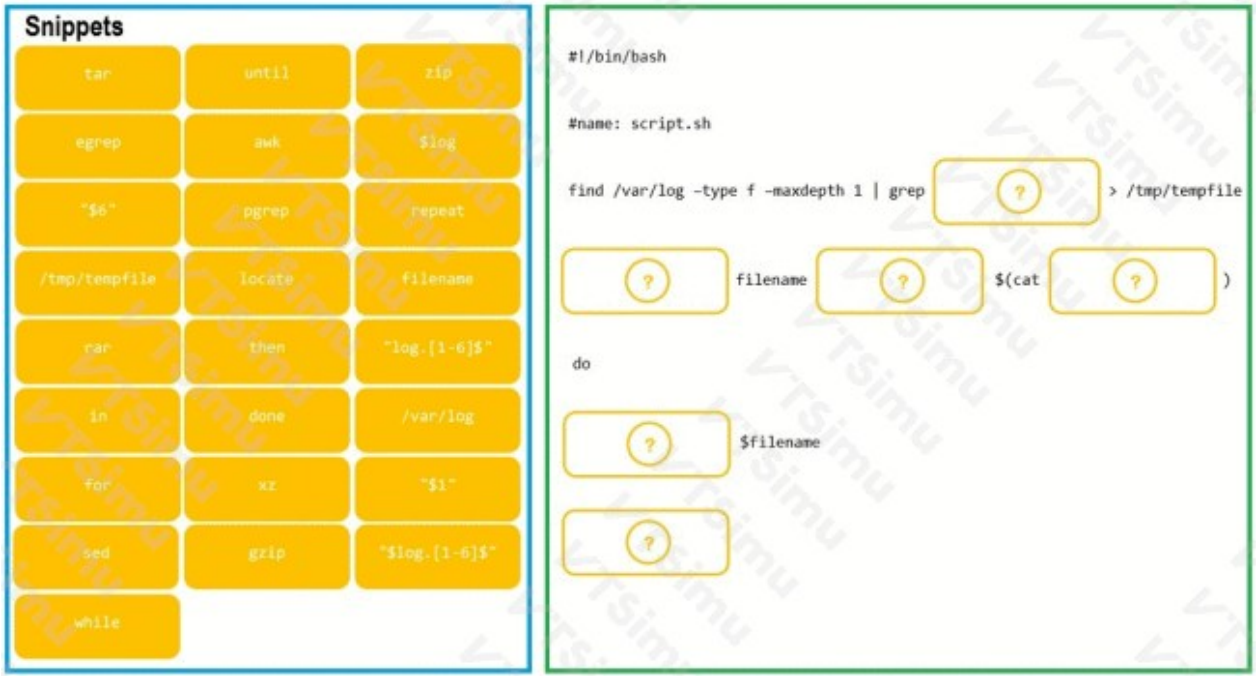

ANSWER:

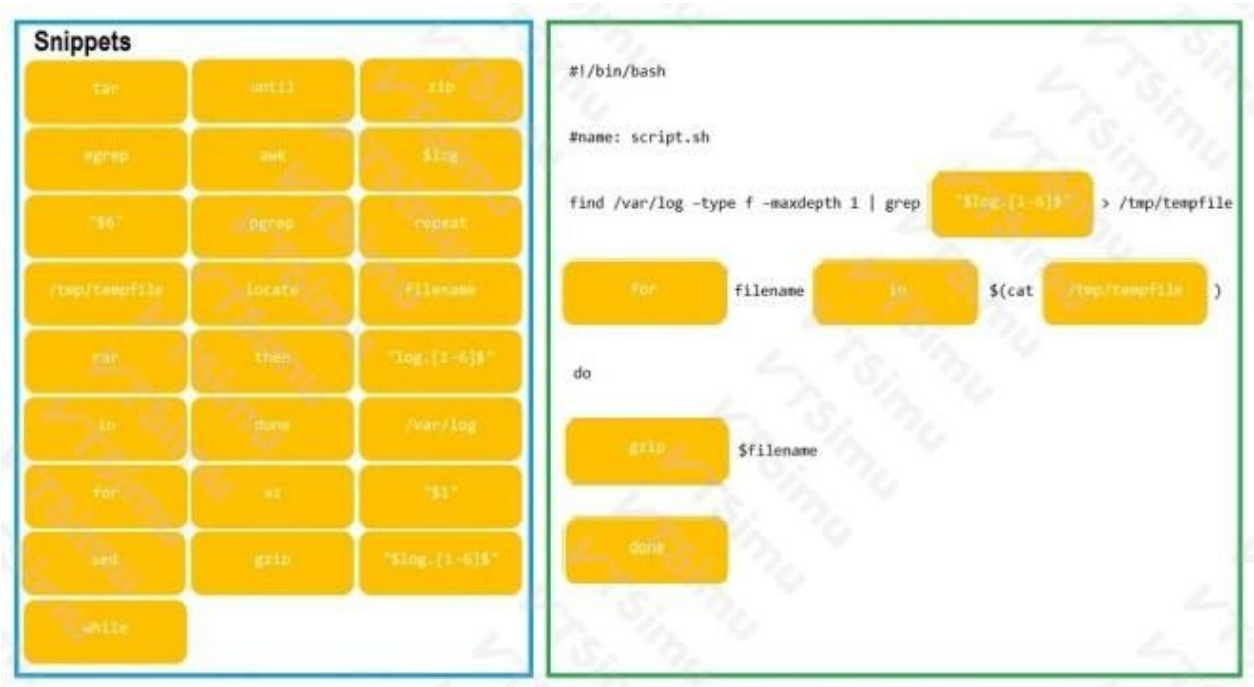

Explanation:

The completed script first uses `find /var/log -type f -maxdepth 1` to list regular files located directly in the `/var/log` directory. The `-maxdepth 1` setting is important because it keeps the search at that directory level instead of descending into subdirectories. The resulting file list is piped to `grep`, where the rotated-log regular expression selects names ending in a log rotation number from 1 through 6. Those matching paths are redirected into `/tmp/tempfile`.

The shell loop then uses `for filename in $(cat /tmp/tempfile)` to process every pathname written by the search pipeline. For each iteration, `gzip $filename` performs the actual compression, and `done` closes the loop. This arrangement separates discovery from processing while ensuring that only the matching rotated logs are passed to the compression utility. GNU Findutils documents `-maxdepth` and file-type tests in the GNU find manual. The GNU Grep manual describes using regular expressions to select matching lines, including end-of-line anchors. Finally, the GNU gzip manual documents gzip as the utility used to compress each selected log file.

QUESTION NO: 18

A Linux administrator is configuring a web service that must read files from the directory `/srv/webcontent` while SELinux is enforcing. Which of the following commands can be used to configure the appropriate persistent SELinux file context and apply it to the directory? (Choose two.)

A. semanage fcontext -a -t httpd_sys_content_t "/srv/webcontent(/.*)?"

B. restorecon -Rv /srv/webcontent

C. chmod -R 777 /srv/webcontent

D. setenforce 0

E. chcon -R -t httpd_sys_content_t /srv/webcontent

ANSWER: A B

Explanation:

`semanage fcontext -a -t httpd_sys_content_t "/srv/webcontent(/.*)?"` adds a persistent SELinux file-context rule. The expression applies the `httpd_sys_content_t` type to the directory and its contents, allowing an

SELinux-confined HTTP service to read the content according to policy. The rule is retained across relabeling and system reboots.

`restorecon -Rv /srv/webcontent` applies the file contexts defined by SELinux policy and local `semanage fcontext` rules to the existing directory tree. The recursive and verbose options cause the directory contents to be processed and report changes. Using these commands together establishes the persistent rule and labels the existing files accordingly. See the semanage-fcontext manual page and the restorecon manual page.

QUESTION NO: 19

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:

Kubernetes is the best fit because it is a container-orchestration platform designed to manage containerized workloads across a cluster of multiple servers (nodes). It provides a control plane that schedules containers onto appropriate nodes, maintains the requested number of application replicas, monitors workload health, and replaces or reschedules failed containers automatically. Kubernetes also supports service discovery, load balancing, rolling updates, configuration management, and scalable deployment patterns. These capabilities allow administrators to operate large, distributed container environments declaratively: they describe the desired application state, and Kubernetes continually works to make the cluster match that state. This is particularly valuable when workloads must be distributed across many hosts while remaining available and able to scale. The Kubernetes documentation describes its role as an open-source platform for managing containerized workloads and services, including automation of deployment, scaling, and management. For an overview of its core orchestration model, see the Kubernetes Concepts Overview. Details on cluster components, including the control plane and worker nodes that enable multi-server orchestration, are available in the Kubernetes Cluster Architecture documentation.

QUESTION NO: 20

A systems administrator needs to install the file installer_0.17-41.2_amd64.deb. Which of the following commands should the administrator use?

A. apt install installer

B. dpkg -i installer_0.17-41.2_amd64.deb

C. rpm -i installer_0.17-41.2_amd64.deb

D. yum localinstall installer_0.17-41.2_amd64.deb

ANSWER: B

Explanation:

The command `dpkg -i installer_0.17-41.2_amd64.deb` is correct because `.deb` is the native binary package format used by Debian and Debian-derived Linux distributions, including Ubuntu. The `dpkg` utility is Debian's low-level package-management tool, and its `-i` (or `--install`) option installs a specified local package archive. The filename identifies a package built for the 64-bit AMD/Intel architecture, as indicated by `amd64`. When run with appropriate administrative privileges, such as through `sudo`, dpkg unpacks the package and registers it in the local package database. A practical command would therefore be `sudo dpkg -i installer_0.17-41.2_amd64.deb`. If the package has

dependencies that are not already installed, an administrator can subsequently use APT to resolve them, commonly with `sudo apt-get -f install`. Debian's dpkg documentation specifies that `--install` installs one or more `.deb` package files, and Ubuntu documentation likewise describes using dpkg to install a local Debian package. See the dpkg manual page and the Ubuntu Debian package documentation.

QUESTION NO: 21

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

```
---
-name: start an instance with a public IP address
 community.abc.ec2_instance:
    name: "public-compute-instance"
    key_name: "comptia-ssh-key"
    vpc_subnet_id: subnet-5cjssh1
    instance_type: instance.type
    security_group: comptia
    network:
       assign_public_ip: true
    image_id: ami-1234568
    tags:
       Environment: Comptia-Items-Writing-Workshop
...
```

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

A. Puppet

B. Git

C. Ansible

D. Terraform

ANSWER: D

Explanation:

Terraform is correct because it is an infrastructure-as-code tool designed to declaratively create, modify, and manage infrastructure resources through cloud-provider APIs. A Terraform configuration can define a cloud virtual machine or instance together with its networking requirements, including a public IP address, network interface, subnet attachment, security rules, and related resources. After the configuration is written, the administrator uses Terraform commands such as `terraform init`, `terraform plan`, and `terraform apply` to initialize the provider, preview the proposed infrastructure changes, and provision the resources. Terraform maintains state so it can associate the declared configuration with the resources that exist in the cloud environment and make controlled updates later. This workflow is especially appropriate when a task involves creating cloud infrastructure rather than only configuring software on an already-created host. The image associated with the question likely depicts Terraform configuration syntax or a Terraform execution workflow, both of which support this identification. HashiCorp's Terraform documentation describes Terraform as an infrastructure-as-code tool for building, changing, and versioning infrastructure safely and efficiently. See the Terraform introduction and the Terraform resource configuration documentation.

QUESTION NO: 22

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

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

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

A. df -i /oracle1

B. fdisk -1 /dev/sdb1

C. lsblk /dev/sdb1

D. du -sh /oracle1

ANSWER: A

Explanation:

`df -i /oracle1` is the appropriate diagnostic command because it reports inode usage for the filesystem that contains `/oracle1`. A filesystem needs both available data blocks and available inodes to create files. Every file and directory consumes an inode, which stores metadata such as ownership, permissions, timestamps, and pointers to the file's data blocks. Consequently, users can be unable to write data even when ordinary space reporting shows free capacity: if the filesystem has exhausted its inode allocation, it cannot create a new file or directory entry.

Using `df` with the `-i` flag displays inode totals, inodes in use, inodes available, and the inode-use percentage for the relevant mounted filesystem. This directly verifies whether inode exhaustion is responsible for the reported write failures. The command can be targeted at the directory path, allowing the system to resolve the correct underlying mount rather than requiring the administrator to first identify its device name. The GNU `df` documentation describes `-i` as listing inode information instead of block usage: GNU Coreutils df invocation. Additional details on Linux filesystem inode usage are available in the df(1) manual page.

QUESTION NO: 23

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

A. docker tag comptia/app:2.1.1 comptia/app:2.1.2

B. docker push comptia/app:2.1.1 comptia/app:2.1.2

C. docker rmi comptia/app:2.1.1 comptia/app:2.1.2

D. docker update comptia/app:2.1.1 comptia/app:2.1.2

ANSWER: A

Explanation:

`docker tag comptia/app:2.1.1 comptia/app:2.1.2` is the correct command because Docker image version labels are tags, and the `docker tag` command applies a new repository-and-tag reference to an existing local image. Here, the existing image identified by `comptia/app:2.1.1` receives the desired `comptia/app:2.1.2` tag. This does not rebuild the application image, alter its filesystem layers, or change its immutable image ID; it creates an additional human-readable reference to the same image. That makes it appropriate when the image was built correctly but labeled with an incorrect version number. After confirming that the new tag is present with a command such as `docker image ls`, the old tag can optionally be removed if it should no longer be used. Docker documents the tag command as creating a tag that refers to a source image, identified by image name, tag, or ID. See the Docker image tag command reference and Docker's image layers documentation for the relationship between tags and the underlying image content.

QUESTION NO: 24

Users connecting to a MySQL database on a Linux system report frequent errors. An administrator reviews the following partial output of the /proc/meminfo and free commands:

total used free

Mem: 1011966 151930 860036

Swap: 1191972 1181716 10256

MemTotal: 1011966 kB

MemFree:

SwapTotal: 1191972 kB

SwapFree: 10256 kB

Which of the following should the administrator do to ensure that the MySQL service will continue to run even if the system runs out of swap space? (Select **two**.)

A. Issue the sysctl -w vm.overcommit_ratio=100 command.

B. Issue the sysctl -w vm.overconunit_memory=2 command. *(Invalid key)*

C. Issue the dd if=/dev/zero of=/swapfile bs=1024 count=65536 command, then run mkswap /swapfile and swapon /swapfile.

D. Issue the mkswap /dev/sda2; swapon -v /dev/sda2 command.

E. Reload the /etc/fstab file by issuing the mount -a command. *(Not sufficient alone)*

F. Use disk quotas on the swap space to prevent MySQL from overutilizing the swap partition.

ANSWER: C D

Explanation:

The displayed memory statistics show that almost all configured swap space is already consumed: only 10,256 kB remains free out of 1,191,972 kB. Increasing available and active swap space gives the kernel additional backing storage for inactive memory pages, reducing the immediate risk that memory allocations needed by MySQL will fail under pressure. Creating a swap file with `dd`, formatting it using `mkswap`, and enabling it with `swapon` is an appropriate way to add swap capacity without repartitioning the disk. The swap file should also be secured with restrictive permissions and added to `/etc/fstab` if it must persist after reboot. Linux swap administration guidance is available in the swapon(8) manual page.

Formatting an available disk partition as swap with `mkswap` and activating it with `swapon` likewise expands the pool of usable swap. This approach is suitable when a dedicated, unused partition such as `/dev/sda2` is available. The mkswap(8) manual page describes creating a Linux swap area before activation. Both methods directly address the dangerously low SwapFree value by supplying additional active swap capacity.

QUESTION NO: 25

A Linux engineer has been notified about the possible deletion of logs from the file /opt/app/logs. The engineer needs to ensure the log file can only be written into without removing previous entries.

```
# lsattz /opt/app/logs
--------------e---  logs
```

Which of the following commands would be BEST to use to accomplish this task?

A. chattr +a /opt/app/logs

B. chattr +d /opt/app/logs

C. chattr +i /opt/app/logs

D. chattr +c /opt/app/logs

ANSWER: A

Explanation:

`chattr +a /opt/app/logs` sets the append-only extended attribute on the specified filesystem object. For a log file, this means processes may write new data only by appending it to the existing end of the file. Existing log entries cannot be overwritten, truncated, or removed while the attribute is active, helping preserve an audit trail and protecting historical records from accidental or unauthorized deletion. This is especially appropriate for logs because normal logging behavior adds successive records rather than editing earlier entries.

The append-only attribute is enforced by the filesystem and is visible with tools such as `lsattr`. Administrative privileges are normally required to set or clear this attribute. If log rotation is used, the administrator must plan for the attribute to be removed temporarily or use a rotation process compatible with append-only handling, because rotating a log commonly involves renaming, truncating, or replacing it. The Linux `chattr` documentation defines `a` as allowing a file to be opened only in append mode and states that it prevents deletion and renaming. See the chattr manual page and the lsattr manual page.

QUESTION NO: 26

A Linux administrator needs to configure a firewall using firewalld. HTTPS traffic must be permitted permanently, and the new configuration must take effect immediately. Which of the following commands should the administrator use? (Choose two.)

A. firewall-cmd --permanent --add-service=https

B. firewall-cmd --reload

C. firewall-cmd --remove-service=https

D. iptables -F

E. systemctl mask firewalld

ANSWER: A B

Explanation:

`firewall-cmd --permanent --add-service=https` adds the HTTPS service to the permanent firewalld configuration for the applicable default zone. The HTTPS service definition normally permits TCP port 443. The `--permanent` option ensures that the rule remains present after firewalld is restarted or the system reboots.

`firewall-cmd --reload` reloads the firewalld configuration and applies permanent configuration changes to the active runtime configuration. Without a reload, a permanent rule does not necessarily affect current traffic immediately. See the firewall-cmd documentation and the firewalld concepts documentation.

QUESTION NO: 27

Which of the following command sequences is used to install Linux software directly from source code?

A. ./config → ./compile → ./copy

B. unzip → gcc → install

C. ./unpack → make exec → make copy

D. ./configure → make → make install

ANSWER: D

Explanation:

The command sequence **./configure → make → make install** is the traditional workflow for building and installing software distributed as source code using the GNU Autotools build system. First, `./configure` examines the host system, checks for required compilers, libraries, and headers, and generates the appropriate build files, commonly a `Makefile`. Configuration options can be supplied at this stage, such as `--prefix=/usr/local`, to select the eventual installation location. Next, `make` reads the generated build instructions and compiles the program and any required components. Finally, `make install` copies the resulting binaries, libraries, documentation, and other installation files to the configured destination. Because writing to system-wide locations normally requires elevated privileges, the final command is often run as `sudo make install`, subject to the organization's software-management policy. This sequence is specifically associated with source packages that include an executable `configure` script; modern projects may instead provide build instructions based on tools such as CMake or Meson. GNU's Automake documentation describes the configure, build, and install stages in this standard process: GNU Automake: Basic Installation. Additional context on GNU build-system conventions is available from the GNU Autoconf manual.

QUESTION NO: 28

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 is currently 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: C D

Explanation:

The `httpd service is currently stopped` statement is supported by the `Active: inactive (dead)` status. This state means systemd is not currently running an Apache HTTP Server process for the `httpd.service` unit. Because Apache is the service expected to listen for and process HTTP requests, its inactive state directly explains why the applications team cannot access the hosted web service. The output establishes the current state, but it does not identify the specific action or event that caused the service to stop.

The `httpd service is not enabled to auto start at boot time` statement is supported by `Loaded: loaded ... disabled`. In systemd, enabling a service creates the appropriate dependency links so that the service is started automatically when the relevant boot target is reached. A disabled unit can still be started manually, but it is not configured for persistent startup across reboots. Enabling and starting the service would address both the immediate availability issue and the lack of automatic startup configuration. Red Hat documents service states and startup management with systemd service management; the systemctl manual also defines enablement and service state behavior.

QUESTION NO: 29

A Linux systems administrator needs to copy files and directories from Server A to Server B. Which of the following commands can be used for this purpose? (Select TWO)

A. rsyslog

B. cp

C. rsync

D. reposync

E. scp

F. ssh

ANSWER: C E

Explanation:

`rsync` and `scp` are both standard tools for transferring files and directories between Linux systems. `rsync` can copy to or from a remote host over SSH using a remote-path form such as `rsync -a /source/ user@server-b:/destination/`. Its archive mode preserves common file attributes and supports efficient subsequent transfers by sending only changed portions of files where possible. It can also recursively synchronize directory trees, making it well suited to routine server-to-server data transfers and backups.

`scp` uses the Secure Shell protocol to securely copy files between hosts. A recursive directory transfer can be performed with a command such as `scp -r /source/ user@server-b:/destination/`. SSH authentication, encryption in transit, and host verification are provided by the underlying SSH connection. Both commands therefore directly satisfy the need to copy files and directories from Server A to Server B, assuming network connectivity, valid credentials, and suitable destination permissions. See the rsync manual page and the OpenSSH scp documentation for syntax and operational details.

QUESTION NO: 30

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

A. git branch -m st

B. git commit -m staging

C. git status -b staging

D. git checkout -b staging

ANSWER: D

Explanation:

The command `git checkout -b staging` is correct because it creates a new local branch named `staging` and immediately checks it out, making it the engineer's active working branch. The `-b` flag combines two normally separate Git operations: creating the branch at the current commit, which is the tip of the currently checked-out `main` branch in this scenario, and switching the working tree and HEAD to the newly created branch. This is useful when a required local branch does not yet exist and the engineer needs to begin work on it right away. Any commits created after the switch will be recorded on `staging`, leaving `main` unchanged unless the branches are later merged. Git documents `git checkout -b <new-branch>` as a convenient form of creating a branch and checking it out in one operation. Modern Git also provides the equivalent command `git switch -c staging`, but the listed `git checkout -b staging` command directly performs the requested task. See the Git checkout documentation and the Pro Git guide to basic branching.

QUESTION NO: 31

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

A. sudo fdisk /dev/sda

B. sudo fdisk -s /dev/sda

C. sudo fdisk -l

D. sudo fdisk -h

ANSWER: C

Explanation:

`sudo fdisk -l` is correct because the `-l` (list) option displays partition tables for the detected block devices. Its output identifies each disk and shows the partitions defined on it, including the partition device path, such as `/dev/sda1` for the first partition on `/dev/sda`. This lets the administrator confirm the name to use in subsequent tasks, such as creating a filesystem, mounting the partition, or adding it to `/etc/fstab`.

When run with elevated privileges, `fdisk -l` can read and report partition-table details comprehensively. The listing includes relevant information such as partition start and end sectors, size, partition type, and boot-related flags where applicable. On systems using GPT partition tables, the command still provides a useful listing of the partition device nodes created for the disk. If a partition was just created, the kernel may need to reread the partition table before the new device node is available; however, `fdisk -l` remains the appropriate command among the choices for verifying the partition's assigned device name. See the fdisk manual page and the Linux kernel partition documentation.

QUESTION NO: 32

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

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

A. Add #!/bin/bash to the bottom of the script.

B. Create a unit file for the new service at /etc/systemd/system/helpme.service.

C. Add #!/bin/bash to the top of the script.

D. Restart the computer to enable the new service.

E. Create a unit file for the new service in /etc/init.d with the name helpme.service in the location.

F. Shut down the computer to enable the new service.

ANSWER: B C

Explanation:

Adding `#!/bin/bash` as the first line of `test.sh` identifies the interpreter that must execute the file. This line, called a shebang, is processed when Linux directly executes an executable script. It must begin with `#!` followed by the valid absolute path of the interpreter. With `/bin/bash` specified, the kernel starts Bash and passes the script to it for interpretation. This is particularly important when the script is launched by a service manager rather than interactively from a shell.

Creating a systemd unit at `/etc/systemd/system/helpme.service` provides the service definition systemd needs to manage the startup task. The unit should define a `[Service]` section whose `ExecStart=` directive uses the absolute path to the executable script, plus appropriate `[Unit]` and `[Install]` sections for its startup relationship and enablement target. After creating or modifying the unit, the administrator should run `systemctl daemon-reload` so systemd reads the new definition, then enable and start it as appropriate. System unit files placed under `/etc/systemd/system` are intended for locally created administrator configuration. See the systemd.service documentation and the GNU Bash Reference Manual.

QUESTION NO: 33

A systems administrator is tasked with setting up key-based SSH authentication. In which of the following locations should the administrator place the public keys for the server?

A. ~/.sshd/authkeys

B. ~/.ssh/keys

C. ~/.ssh/authorized_keys

D. ~/.ssh/keyauth

ANSWER: C

Explanation:

The correct location is `~/.ssh/authorized_keys`. For standard OpenSSH public-key authentication, the public key belonging to a client/user is installed on the SSH server in the target account's `authorized_keys` file. When that user attempts to connect, the SSH server uses this file to determine whether the public key presented during authentication is authorized for that account. The tilde represents the home directory of the account that will log in; for example, the file for user `admin` would normally be `/home/admin/.ssh/authorized_keys`. For the root account, it is commonly `/root/.ssh/authorized_keys`.

The directory and file must also be protected with suitable ownership and permissions because OpenSSH checks them as part of its security model. A typical configuration is a user-owned `~/.ssh` directory with mode `700` and a user-owned `authorized_keys` file with mode `600`. The server's `sshd_config` setting `AuthorizedKeysFile` specifies this lookup path and defaults to a path based on `.ssh/authorized_keys`. See the OpenSSH sshd_config AuthorizedKeysFile documentation and the sshd manual.

QUESTION NO: 34

A Linux administrator is changing the default system umask. The newly created files should have the -rw-r----- permission, and the newly created directories should have the drwxr----- permission. Which of the following commands should the administrator use?

A. umask 0035

B. umask 0036

C. umask 0037

D. umask 0038

E. umask 0027

ANSWER: E

Explanation:

`umask 0027` is correct because a umask specifies which permission bits are removed from the base creation modes used by applications. New regular files conventionally begin with a maximum mode of `0666`, rather than `0777`, so that execute permission is not automatically assigned. New directories begin with a maximum mode of `0777`. The mask `0027` removes write permission from the group class and all permissions from the other class. Applying that mask to `0666` results in `0640`, displayed as `-rw-r-----`: the owner can read and write, the group can read, and others have no access. Applying the same mask to `0777` results in `0750`, displayed as `drwxr-----`: the owner has full access, the group can read and traverse the directory, and others have no access. The leading zero simply makes the intended octal representation explicit; `umask 027` is equivalent in shells that accept octal masks. This behavior is defined by the system umask interface and is commonly used to establish secure default access for newly created content. See the Linux umask(2) manual page and the POSIX umask specification.

QUESTION NO: 35

In order to copy data from another VLAN, a systems administrator wants to temporarily assign IP address 10.0.6 5/24 to the newly added network interface enp1s0f1. Which of the following commands should the administrator run to achieve the goal?

A. ip addr add 10.0.6.5/24 dev enp1s0f1

B. echo "IPV4_ADDRESS=10.0.6.5/24" > /etc/sysconfig/network-scripts/ifcfg-enplsOfl

C. ifconfig 10.0.6.5/24 enpsls0f1

D. nmcli conn add lpv4.address-10.0.6.5/24 ifname enpls0f1

ANSWER: A

Explanation:

`ip addr add 10.0.6.5/24 dev enp1s0f1` is the appropriate command because it uses the iproute2 `ip address` utility to add the IPv4 address and its /24 prefix directly to the specified interface. The prefix length defines the subnet mask, so /24 configures the address with a 255.255.255.0 mask. This is suitable for a temporary address assignment needed to access a host or transfer data on another VLAN, because the address is applied immediately to the running network configuration and is not inherently saved as a persistent connection profile. The command explicitly identifies both the address to assign and the target device, which is important on systems with multiple network adapters. The interface name must exactly match the newly added device, so the command uses `enp1s0f1`. The `ip` command is the current standard Linux networking tool for viewing and modifying addresses, routes, and links. The Linux iproute2 manual documents the `ip address add` syntax and behavior at ip-address(8). Red Hat also documents temporary IP address configuration using the `ip addr add` command at Configuring and managing networking.

QUESTION NO: 36

A Linux administrator is implementing a stateful firewall on the Linux server. Which of the following iptables options will be required to build the stateful rules? (Select two).

A. --name established

B. -m recent

C. -m conntrack

D. --state

E. --remove

F. -j DROP

ANSWER: C D

Explanation:

`-m conntrack` is used to load the conntrack match extension, allowing iptables rules to evaluate information maintained by the kernel connection-tracking subsystem. This is the core mechanism that lets a firewall recognize packets as belonging to an existing flow, a related flow, or a new connection instead of treating every packet independently. Connection tracking supports common stateful policies, such as permitting reply traffic for connections that were initiated from the protected host or network.

`--state` specifies the connection state that a rule should match when using iptables' legacy state match functionality. Administrators can use it to select states such as `NEW`, `ESTABLISHED`, `RELATED`, and `INVALID`. In current iptables usage, the conntrack match generally uses the closely related `--ctstate` parameter, for example `-m conntrack --ctstate ESTABLISHED,RELATED`; however, `--state` represents the state-based matching concept requested by the question. Together, connection tracking and connection-state matching provide the information needed to construct stateful filtering rules. See the iptables-extensions manual and the Netfilter iptables documentation.

QUESTION NO: 37

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

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

A. gpg /dev/sdcl

B. pvcreate /dev/sdc

C. mkfs.ext4 /dev/mapper/LUKS0001 -L ENCRYPTED

D. umount /dev/sdc

E. fdisk /dev/sdc

F. mkfs.vfat /dev/mapper/LUKS0001 -L ENCRYPTED

G. wipefs -a /dev/sdb1

H. cryptsetup luksFormat /dev/sdc1

ANSWER: C D H

Explanation:

To create an encrypted ext4 volume with LUKS, the existing mounted block device must first be unmounted so that its contents and metadata are not being modified while encryption setup begins. The administrator should then initialize the target partition as a LUKS container with `cryptsetup luksFormat /dev/sdc1`. This writes the LUKS header and configures the passphrase-based encryption container. After the LUKS container has been opened and mapped as `/dev/mapper/LUKS0001`, an ext4 filesystem can be created on that decrypted mapper device using `mkfs.ext4 /dev/mapper/LUKS0001 -L ENCRYPTED`. The filesystem must be placed on the mapped device rather than directly on the LUKS partition, because the mapper device provides the decrypted block-device view used by the filesystem layer. The `-L ENCRYPTED` argument assigns a filesystem label, which can make later identification and mounting easier. These steps follow the standard LUKS workflow of unmounting the source, formatting a LUKS container, opening it, and formatting the resulting mapped device. See the cryptsetup-luksFormat manual and the mkfs.ext4 manual.

QUESTION NO: 38

A Linux administrator needs to investigate why the `nginx.service` unit failed during the current boot. Which of the following commands will show journal entries for that unit from the current boot only? (Choose two.)

A. journalctl -u nginx.service -b

B. journalctl -b -u nginx.service

C. journalctl --vacuum-time=1d

D. systemctl enable nginx.service

E. tail -f /etc/nginx/nginx.conf

ANSWER: A B

Explanation:

`journalctl -u nginx.service -b` filters journal records to the specified systemd unit and restricts the results to the current boot. This provides the service messages needed to diagnose startup errors while excluding records from earlier boots.

`journalctl -b -u nginx.service` uses the same journalctl filters in a different valid order. Both commands request records for `nginx.service` from boot ID zero, which represents the current boot. See the journalctl documentation.

QUESTION NO: 39

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

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

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

A. Add #!/bin/bash to the bottom of the script.

B. Create a unit file for the new service in /etc/systemd/system/ named helpme.service.

C. Add #!/bin/bash to the top of the script.

D. Restart the computer to enable the new service.

E. Create a unit file for the new service in /etc/init.d with the name helpme.service in the location.

F. Shut down the computer to enable the new service.

ANSWER: B C

Explanation:

Adding `#!/bin/bash` as the first line supplies the script's interpreter directive, commonly called a shebang. When an executable text file is invoked directly, the operating system uses this first line to identify the program that must interpret its contents. With Bash specified at the top, executing `test.sh` runs its commands through the Bash shell. The original double-slash form, `#!//bin/bash`, resolves to the same path on Linux, but the conventional and clear form is `#!/bin/bash`.

To have the task managed and started by systemd, the administrator should create a service unit such as `/etc/systemd/system/helpme.service`. Its `[Service]` section should use an appropriate absolute `ExecStart=` path for the script, and its `[Install]` section can specify a boot target such as `multi-user.target`. After creating or

changing the unit, systemd must reload its unit definitions, and the service can be enabled for startup with `systemctl enable helpme.service`. These steps establish both a valid executable script and a systemd-managed startup service. See the systemd.service documentation and the GNU Bash Reference Manual.

QUESTION NO: 40

A Linux administrator needs to troubleshoot a system that becomes unresponsive during periods of heavy disk activity. Which of the following commands would be MOST helpful for identifying disk I/O pressure? (Choose two.)

A. iostat -xz 1

B. iotop -o

C. free -h

D. ip addr show

E. uname -r

ANSWER: A B

Explanation:

`iostat -xz 1` reports CPU and device I/O statistics repeatedly at one-second intervals. The extended report includes values such as device utilization, average queue size, await time, and read/write throughput. These metrics help identify devices that are saturated or experiencing elevated I/O latency.

`iotop -o` displays processes and threads that are actively performing I/O. The `-o` option limits the display to tasks with current I/O activity, helping an administrator associate disk pressure with the responsible workload. Together, these commands provide both device-level and process-level evidence. The iostat(1) manual page and the iotop(8) manual page document these utilities.

QUESTION NO: 41

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

A. Add the line DenyUsers root to the /etc/hosts.deny file.

B. Set PermitRootLogin to no in the /etc/ssh/sshd_config file.

C. Add the line account required pam_nologin. so to the /etc/pam.d/sshd file.

D. Set PubKeyAuthentication to no in the /etc/ssh/ssh_config file.

ANSWER: B

Explanation:

Set `PermitRootLogin` to `no` in the server-side `/etc/ssh/sshd_config` configuration file. The administrative account on a Linux system is ordinarily `root`, and this OpenSSH daemon directive explicitly prevents that account from authenticating through SSH. This is an important hardening control because it removes direct remote root access while leaving normal SSH access available for named, nonprivileged administrator accounts. Administrators can then sign in with an individual account and use controlled privilege elevation, such as `sudo`, for administrative tasks. That approach improves accountability because authentication and command activity can be associated with a specific user rather than a shared root login. After changing the daemon configuration, the administrator should validate it with `sshd -t` before reloading or restarting the SSH service, helping avoid an accidental configuration error that could block remote management. OpenSSH documents that `PermitRootLogin no` prohibits root from logging in via SSH. See the OpenSSH sshd_config PermitRootLogin documentation and the Red Hat OpenSSH security guidance.

QUESTION NO: 42

A Linux administrator is alerted to a storage capacity issue on a server without a specific mount point or directory. Which of the following commands would be MOST helpful for troubleshooting? (Choose two.)

A. parted

B. df

C. mount

D. du

E. fdisk

F. dd

G. ls

ANSWER: B D

Explanation:

`df` is appropriate because it reports filesystem disk-space usage, including total capacity, used space, available space, usage percentage, and associated mount points. Running `df -h` presents this information in human-readable units and helps the administrator quickly identify which mounted filesystem is approaching or has reached capacity. The GNU coreutils documentation describes `df` as reporting file-system disk-space usage: GNU Coreutils: df invocation.

`du` complements this by estimating space consumed by files and directories. After `df` identifies a full filesystem, commands such as `du -xhd1 /mountpoint` can identify the largest top-level directories without crossing onto other filesystems. This provides the directory-level detail needed to locate the data responsible for the capacity alert. The GNU Coreutils manual documents that `du` reports the amount of disk space used by files and directories: GNU Coreutils: du invocation. Together, these commands establish both where the capacity problem exists and what data is consuming the affected filesystem.

QUESTION NO: 43

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

A)

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

B)

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

C)

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

D)

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

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: D

Explanation:

The command sequence shown in *Option D* is correct because creating and mounting a new XFS filesystem requires the operations to occur in the proper order: make the target mount-point directory, create the filesystem with the XFS formatting utility, and mount the formatted block device at `/ops/app`. A 1 GB filesystem is established by allocating or preparing a 1 GB backing device or logical volume before formatting it. The `mkfs.xfs` command writes the XFS filesystem metadata to that device, making it usable as an XFS filesystem. The `mount` command then attaches that filesystem to the existing `/ops/app` directory, where its contents become accessible.

Because the requirement states that the filesystem should be mounted temporarily, the mount should be performed interactively and should not require a persistent `/etc/fstab` entry. A temporary mount remains active until it is explicitly unmounted or the system is rebooted. This workflow follows standard Linux filesystem administration practice: prepare storage, create the filesystem, ensure the mount point exists, and mount the filesystem. See the mkfs.xfs manual page and the mount manual page for command behavior and syntax.

QUESTION NO: 44

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

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

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

A)

```
IPtables -A FORWARD -i eth0 -p tcp --dport 80 -j ACCEPT
IPtables -A FORWARD -i eth0 -p tcp --dport 22 -j ACCEPT
```

B)

```
IPtables -A INPUT -i eth0 -p tcp --dport 80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --dport 22 -j ACCEPT
```

C)

```
IPtables -A INPUT -i eth0 -p tcp --sport 80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --sport 22 -j ACCEPT
```

D)

```
IPtables -A INPUT -i eth0 -p tcp --dport :80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --dport :22 -j ACCEPT
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: B

Explanation:

Option B is correct because it applies the firewall remediation needed to restore access to services that are being blocked by the packet-filtering rules identified in the kernel log. When a host is unreachable while kernel messages indicate firewall processing or rule-related drops, the administrator must correct the active filtering policy so that legitimate traffic can reach the affected service. The remediation must be applied to the appropriate traffic path and service port, while preserving a secure default policy rather than broadly disabling host protection. After applying the rule change, the administrator should validate that the web service is listening, confirm that the firewall permits the intended protocol and destination port, and test connectivity from an external client. The change should also be made persistent using the firewall-management method used by the distribution so it survives a reboot. CompTIA Linux+ objectives include configuring and troubleshooting network security controls, including host-based firewall rules. The netfilter/iptables documentation further describes how packet-filter rules are evaluated and managed: CompTIA exam objectives and iptables manual page.

QUESTION NO: 45

A Linux systems administrator needs to verify that a disk image downloaded from an internal repository has not been modified. The repository provides a file named `server.img.sha256`. Which of the following commands can be used to verify the image against the provided SHA-256 checksum? (Choose two.)

A. sha256sum -c server.img.sha256

B. sha256sum server.img

C. chmod 600 server.img

D. rpm -V server.img

E. tar -tzf server.img

ANSWER: A B

Explanation:

`sha256sum -c server.img.sha256` reads the checksum file, calculates the SHA-256 digest of each referenced file, and reports whether the calculated value matches the expected value. This is the standard verification form when the checksum file is in the format generated by `sha256sum`.

`sha256sum server.img` calculates and displays the SHA-256 digest of the image. An administrator can compare this output with the expected digest in `server.img.sha256`. Although the `-c` form is less error-prone because it performs the comparison automatically, both commands can be used for verification. See the GNU Coreutils SHA-2 utilities documentation.

QUESTION NO: 46

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

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

A. systemctl unmask mariadb

B. journalctl —g mariadb

C. dnf reinstall mariadb

D. systemctl start mariadb

E. chkconfig mariadb on

F. service mariadb reload

ANSWER: A D

Explanation:

`systemctl unmask mariadb` is required when the MariaDB unit has been masked. In systemd, masking is a stronger restriction than simply disabling a service: it prevents activation of the unit entirely, including manual starts and starts triggered by dependencies. The unmask operation removes the mask, restoring systemd's ability to load the service unit file and process start requests for MariaDB.

After the unit is unmasked, `systemctl start mariadb` starts the database service immediately for the current boot. This sequence directly addresses a service that cannot start because its unit is masked: first remove the administrative block, then request normal service activation. Administrators can subsequently verify its state with `systemctl status mariadb` and, if the database should start automatically after future reboots, separately enable it with the appropriate `systemctl enable` command. The systemctl documentation describes masking as making a unit impossible to start and documents `unmask` as reversing that state. See the systemctl manual and the systemd unit documentation.

QUESTION NO: 47

A Linux administrator created a virtual clone of a physical server and would like to remove any existing entries related to SSH keys from outside entities on the virtual clone. Which of the following files should the administrator remove? (Select two).

A. ~/.ssh/authorized_keys

B. ~/.ssh/known_hosts

C. /etc/ssh/ssh_config

D. ~/.ssh/config

E. /etc/ssh/sshd_config

F. /etc/ssh/ssh_host_rsa_key.pub

ANSWER: A B

Explanation:

The correct files to remove are `~/.ssh/authorized_keys` and `~/.ssh/known_hosts`. The `authorized_keys` file holds public keys belonging to remote users or systems that have been authorized to authenticate to the local account. Because a virtual clone inherits the source machine's filesystem, retaining this file could allow entities that were trusted on the original server to continue accessing the cloned account with their existing SSH keys. Removing the inherited file eliminates those prior key-based authorization entries; new access can then be explicitly configured for the clone as needed.

The `known_hosts` file stores SSH host-key entries for remote hosts previously contacted by that user. These are identities of outside SSH servers that the original machine had recorded and trusted. Removing inherited entries gives the cloned system a clean SSH client trust database, so connections to external hosts are evaluated and recorded in the context of the new virtual instance. OpenSSH documents the purpose and location of both user-level files in ssh_config(5). The server-side use of authorized public-key files is also documented in sshd(8).

QUESTION NO: 48

NO: 22:

An organization's business office needs to collaborate on quarterly reports and asks the systems administrator to create a finance group. The administrator creates the following directory:

drwxrwxr-x Business Business 4096 Oct 15 16:40 Q1

Which of the following commands will allow all members of the finance group to have group ownership of all files created in the Q1 directory while ensuring that members of the business group preserve access? (Select two).

A. chgrp Finance Q1/

B. chmod g+s Q1/

C. chmod u+s Q1/

D. chmod a+x Q1/

E. chown Finance Q1/

F. chmod g+x Q1/

ANSWER: A B

Explanation:

`chgrp Finance Q1/` is required because it changes the group ownership of the `Q1` directory to the Finance group. The directory's existing mode grants group read, write, and execute permissions, so Finance group members can collaborate in the directory once its group is changed. The existing permissions for "others" are read and execute, which retain directory traversal and read access for users who are not in the Finance group, including Business users where applicable.

`chmod g+s Q1/` sets the set-group-ID (setgid) bit on the directory. On Linux, files and subdirectories subsequently created within a directory with setgid enabled inherit that directory's group ownership rather than receiving the creating user's primary group. Therefore, after the directory group has been set to Finance, newly created quarterly-report files will be

owned by the Finance group. This provides the intended shared-group ownership model without requiring an administrator to manually change each new file's group.

The GNU Coreutils documentation describes changing a file's group with `chgrp`: GNU chgrp invocation. The directory setgid inheritance behavior is documented in the Linux `chmod` manual page: chmod(1).

QUESTION NO: 49

A Linux administrator needs to configure the directory `/srv/engineering` so that newly created files inherit the `engineering` group and members of that group can create files in the directory. Which of the following commands should the administrator use? (Choose two.)

A. chown :engineering /srv/engineering

B. chmod 2770 /srv/engineering

C. chmod 1770 /srv/engineering

D. chgrp -R root /srv/engineering

E. chmod 4700 /srv/engineering

ANSWER: A B

Explanation:

`chown :engineering /srv/engineering` changes the owning group of the directory to `engineering`. This establishes the group that will be inherited when the directory set-group-ID bit is enabled.

`chmod 2770 /srv/engineering` sets the set-group-ID bit and assigns `rwx` permissions to the owner and group while denying access to others. On a directory, the set-group-ID bit causes newly created files and subdirectories to inherit the directory's group ownership. The group write and execute permissions allow group members to create and access directory entries. See the chown(1) manual page and the GNU Coreutils documentation for set-group-ID directories.

QUESTION NO: 50

A Linux administrator needs to identify all regular files under `/opt/app` that have no owning user in the local account database. Which of the following commands will accomplish this task?

A. find /opt/app -type f -nouser

B. find /opt/app -type f -user root

C. find /opt/app -type f -perm 000

D. find /opt/app -type d -nogroup

ANSWER: A

Explanation:

`find /opt/app -type f -nouser` searches the specified directory hierarchy for regular files whose numeric user IDs cannot be mapped to an account in the system's configured user database. The `-type f` predicate limits results to regular files, and `-nouser` selects files owned by nonexistent users.

This condition can occur after an account is removed while files owned by that account remain on a local filesystem. Before changing ownership or deleting returned files, the administrator should verify that the accounts are genuinely obsolete and review the results. The find(1) manual page documents the `-nouser` and `-type` predicates.

QUESTION NO: 51

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

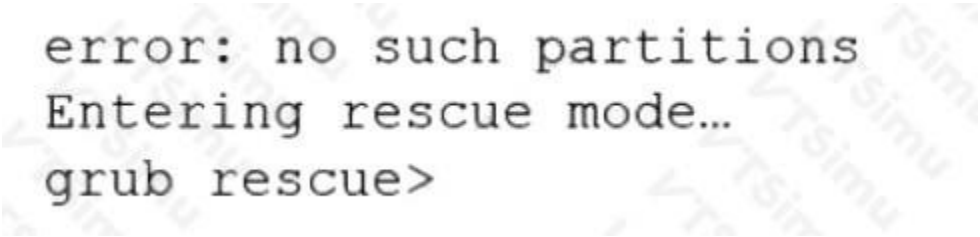

Which of the following actions will resolve this issue? (Choose two.)

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. Correct the partition reference in /etc/default/grub, regenerate the GRUB configuration, 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: B D

Explanation:

Executing `grub-install /dev/sdX` reinstalls the GRUB bootloader to the disk's boot area. The target must be the whole boot disk, such as `/dev/sda`, rather than an individual partition such as `/dev/sda1`. This restores the bootloader components required to begin loading GRUB's configuration and the selected operating system. In a recovery scenario, this command is normally run after booting suitable recovery media and mounting or chrooting into the installed Linux system. GNU documents that `grub-install` installs GRUB to the specified device; see the GNU GRUB installation documentation.

Correcting the partition reference in `/etc/default/grub`, regenerating the GRUB configuration, and rebooting ensures that GRUB uses the correct boot or root location. Settings in this file are input to the generated `grub.cfg`; after making the correction, an administrator should generate the active configuration with the distribution's supported command, commonly `update-grub` or `grub-mkconfig -o /boot/grub/grub.cfg`. This makes the repaired partition setting available at the next boot. GNU's GRUB simple configuration documentation describes how `/etc/default/grub` is used when creating the configuration.

QUESTION NO: 52

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

A. wget

B. ssh-keygen

C. ssh-keyscan

D. ssh-copy-id

E. ftpd

F. scp

ANSWER: D F

Explanation:

`ssh-copy-id` is specifically designed to install a local public SSH key on a remote host. It connects using SSH, commonly authenticating with the administrator's existing password or another available identity, and appends the selected public key to the remote account's `~/.ssh/authorized_keys` file. This is the normal, purpose-built method for enabling public-key authentication on managed Linux servers. The OpenSSH manual documents that it uses SSH to log in to a remote machine and install a public key in the remote user's authorized-keys file: OpenSSH ssh-copy-id manual.

`scp` can also copy a public-key file to a remote server because it securely transfers files over an SSH connection. For example, an administrator can copy a `.pub` file to the remote user's home directory, then log in and place its contents into `~/.ssh/authorized_keys` with suitable ownership and permissions. Although `scp` does not automatically install the key as an authorized identity, it is a valid command for copying the key file itself to the remote system. The OpenSSH documentation describes `scp` as a secure file-copy program that transfers files between hosts over SSH: OpenSSH scp manual.