

DUMPS ARENA

Linux Foundation Certified IT Associate

Linux Foundation LFCA

Version Demo

Total Demo Questions: 7

Total Premium Questions: 79

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Linux Fundamentals	12
Topic 2, System Administration Fundamentals	19
Topic 3, Cloud Computing Fundamentals	20
Topic 4, Security Fundamentals	17
Topic 5, DevOps Fundamentals	11
Total	79

QUESTION NO: 1

A directory is owned by `root:analysts` and has permissions `dr--r-----`. A member of the analysts group knows the name of a readable file inside the directory but cannot open it. Which permission is missing from the directory?

- A. Read permission
- B. Execute permission
- C. Write permission
- D. Set-group-ID permission

ANSWER: B

Explanation:

The execute permission is missing from the directory. For a directory, execute permission permits traversal: a user needs it to access files within the directory, even when the user already knows a file name and the file itself is readable. Group read permission allows analysts to list directory entries, but it does not allow them to traverse the directory and access those entries. Write permission would be needed to create, rename, or remove directory entries, not merely to open an existing file.

QUESTION NO: 2

An IT associate is working with a video streaming provider. What is a common approach to reduce bandwidth cost while supporting the same or an increasing number of video streams and customers?

- A. Stop serving the video streams when a certain price limit is reached.
- B. Reduce the bitrate to reduce the size of the video streams.
- C. Serve the video streams only to a limited number of users at a time.
- D. Use a content delivery network to distribute the streams.

ANSWER: D

Explanation:

Use a content delivery network to distribute the streams. A CDN places cached copies of content, or delivery endpoints, at geographically distributed edge locations close to viewers. Instead of every viewer retrieving video data directly from a single origin infrastructure, requests can be served by a nearby edge location. This reduces repeated long-distance transfers from the origin, lowers origin-network egress and infrastructure load, and improves scalability as the audience grows. For video streaming, CDNs are especially useful because popular segments can be cached and delivered repeatedly to many users without requiring the origin server to transmit each copy independently. A CDN also provides a distributed delivery architecture that can absorb traffic spikes and reduce latency, helping maintain a consistent viewing experience while serving more concurrent customers. The Linux Foundation's introductory cloud material identifies CDNs as a cloud service used to deliver content efficiently, while AWS describes how CDN edge locations cache and serve content closer to users. See the [Linux Foundation Introduction to Cloud Infrastructure Technologies](#) and [AWS CloudFront Developer Guide](#).

QUESTION NO: 3

A web service managed by `systemd` fails shortly after startup. Which command is the most appropriate to view log messages recorded specifically for the service named `webapp.service` during troubleshooting?

- A. `journalctl -u webapp.service`
- B. `dmesg webapp.service`
- C. `tail -f webapp.service`

D. `ps -ef webapp.service`

ANSWER: A

Explanation:

`journalctl -u webapp.service` is correct because the `-u` option filters the systemd journal by unit name. This lets an administrator view messages emitted by the service and by systemd while starting, stopping, or restarting that unit. `systemctl status` can show a summary and recent messages, but `journalctl -u` is better for reviewing the unit's journal entries in detail. `dmesg` is primarily for kernel messages, while `tail -f` requires knowing a particular log file and does not inherently filter by systemd unit.

QUESTION NO: 4

Which of the following is a commonly used open source software used to connect to SSL/TLS VPN services?

- A. OpenVPN
- B. GnuVPN
- C. NordVPN
- D. VPNConnect

ANSWER: A

Explanation:

OpenVPN is a widely used open-source VPN solution designed to establish secure remote-access and site-to-site VPN connections using SSL/TLS for authentication, key exchange, and encryption. Its client and server software are available for Linux and many other operating systems, making it a common choice in cross-platform enterprise and personal VPN deployments. OpenVPN uses the TLS protocol to create and manage a secure control channel, while encrypted data traffic is carried through an established VPN tunnel. It supports certificate-based authentication as well as other authentication methods, helping administrators securely verify users and systems connecting to a VPN service. The project is distributed under an open-source license, and its software can be deployed and configured independently rather than being restricted to a single commercial VPN provider. The OpenVPN community documentation describes it as an open-source VPN solution and provides installation and configuration resources for its client and server components. For further technical details, see the [OpenVPN Community Resources](#) and the [OpenVPN Community Documentation](#).

QUESTION NO: 5

A developer has edited several tracked files in a Git working directory but has not run `git add`. Which command shows the line-by-line changes that have not yet been staged?

- A. `git diff`
- B. `git diff --staged`
- C. `git log`
- D. `git status --branch`

ANSWER: A

Explanation:

`git diff` displays changes between the working tree and the staging area, which means it shows modifications that have not yet been staged with `git add`. This is useful for reviewing edits before deciding which changes to stage. `git diff --`

`staged` instead compares staged changes with the most recent commit. `git status` identifies modified files but does not normally display the detailed line-by-line differences.

QUESTION NO: 6

A company's IT associate has been asked to switch to single-user mode on a running Linux server in order to perform some troubleshooting. Of the options below, which is the best way of doing this?

- A. `su single`
- B. `init singleuser`
- C. `lsu 1`
- D. `init 1`

ANSWER: D

Explanation:

`init 1` is the traditional command for changing a running Linux system to runlevel 1, commonly called single-user mode. Runlevel 1 starts a minimal maintenance environment intended for administrative troubleshooting and recovery tasks, rather than the normal multiuser services and login environment. The `init` command communicates the requested runlevel change to the system's `init` process, so it performs a system-wide state transition rather than merely opening a different shell or changing the current user context.

On modern distributions that use `systemd`, traditional runlevel commands are retained for compatibility and are mapped to corresponding `systemd` targets. In this model, runlevel 1 maps to `rescue.target`; administrators may also use `systemctl isolate rescue.target` when `systemd`-specific administration is appropriate. The command in the answer remains the best choice among the supplied commands because it uses the established runlevel syntax for entering single-user maintenance mode on a running system. Administrative access is required, and remote administrators should take care because services such as networking may be stopped during the transition. See the [systemd special targets documentation](#) and the [systemctl documentation](#).

QUESTION NO: 7

A developer has created and committed a change on the local `main` branch. The corresponding remote repository is configured as `origin`. Which command publishes the new local commit to the remote `main` branch?

- A. `git pull origin main`
- B. `git fetch origin main`
- C. `git push origin main`
- D. `git clone origin main`

ANSWER: C

Explanation:

`git push origin main` is correct because it sends commits from the local `main` branch to the `main` branch in the remote repository named `origin`. A commit records the change only in the local repository until it is pushed. `git pull` retrieves and integrates remote changes, while `git fetch` retrieves remote references without publishing local commits. `git clone` creates a new local copy of an existing repository.