

Linux Foundation LFCT

Linux Foundation Certified Cloud Technician

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/lfct>

Latest Version: 6.0

Question: 1

A network administrator suspects that a firewall is blocking incoming connections to a web server running on port 443. Which command should be used to verify if the firewall is indeed allowing HTTPS traffic on port 443?

- A. `iptables -L | grep 443`
- B. `nmap -p 443 webserver_ip`
- C. `curl https://webserver_ip`
- D. `ss -tuln | grep :443`

Answer: A

Explanation:

`iptables -L | grep 443` checks the current iptables rules for any that specifically mention port 443, which is the default port for HTTPS traffic. This command is most directly related to diagnosing firewall configurations and determining whether the rules are set to allow or deny traffic on port 443. Option B is incorrect. `nmap -p 443 webserver_ip` scans the specified port to see if it's open, which can indicate whether external access is possible but doesn't specify if iptables rules are the cause of any blockage. Option C is incorrect. `curl https://webserver_ip` attempts to access the web server via HTTPS, testing if the service is reachable. While failure might suggest a block, it doesn't confirm the firewall settings directly. Option D is incorrect. `ss -tuln | grep :443` shows if any service is listening on port 443 but doesn't indicate whether firewall rules are permitting or blocking incoming connections on that port.

Question: 2

A DevOps team is managing a fleet of servers across multiple data centers using a GitOps workflow. They want to ensure that any script executed for maintenance tasks is done so consistently and that the output from each server is captured for audit purposes. Which tool or approach would best fit this requirement within a GitOps framework?

- A. Manual execution of scripts via SSH and manual logging of the output on a shared document
- B. Use of a centralized configuration management tool like Ansible, with scripts and their execution policies stored in a Git repository and outputs logged to a centralized logging system
- C. Sending scripts via email to data center administrators and asking them to run the scripts and reply with the output
- D. Utilizing custom bash scripts that are manually copied to each server and executed, with outputs emailed back to the DevOps team

Answer: B

Explanation:

This approach aligns with GitOps principles by using Git as the single source of truth for script storage and execution policies, ensuring consistent execution across all servers. Ansible's ability to log outputs to a centralized system further supports auditability and compliance with GitOps practices. Option A is incorrect. Manual execution and logging are prone to human error and inconsistency, which contradicts the GitOps emphasis on automation and repeatability. Option C is incorrect. Relying on email for script distribution and execution results collection is inefficient, lacks traceability, and does not support the automated, version-controlled approach of GitOps. Option D is incorrect. Custom bash scripts manually copied and executed do not provide the consistency, automation, or auditability required in a GitOps workflow, making this approach less suitable.

Question: 3

You discover that a user account named 'archive' should no longer have login access to the system. Which command can you use to disable login for 'archive' without deleting the account or its files?

- A. userdel archive
- B. passwd -l archive
- C. usermod -L archive
- D. chage -E 0 archive

Answer: B

Explanation:

passwd -l archive locks the password for the user 'archive', effectively disabling login without removing the account or its data. This is a common way to disable an account temporarily or permanently without deletion. Option A is incorrect. userdel is used to delete a user account from the system, which also removes the account's home directory if not instructed otherwise, not just disable login capabilities. Option C is incorrect. While usermod -L can lock the user's password, similar to passwd -l, making it an alternative way to disable login, passwd -l is more directly associated with the task of locking an account for login purposes. Option D is incorrect. chage -E 0 archive sets the account's expiration date to day 0, which effectively disables the account but is a more drastic measure compared to simply locking the password. It's used for different purposes, like completely expiring the account, not just disabling logins.

Question: 4

After changing the system's timezone to 'America/New_York', a Linux system administrator needs to verify the current system timezone setting. Which command can they use to check the timezone?

- A. timedatectl
- B. cat /etc/timezone
- C. date +%Z
- D. ls /usr/share/zoneinfo/America/New_York

Answer: A

Explanation:

timedatectl without any additional arguments displays detailed information about the current system time and timezone. It's a part of systemd and provides a straightforward method to check the timezone on systems using systemd. Option B is incorrect. While cat /etc/timezone might show the timezone on some Debian-based systems, it's not universally applicable across all Linux distributions, especially on systems that do not use this file to store the timezone setting. Option C is incorrect. date +%Z displays the timezone abbreviation, which can confirm the timezone setting but does not provide as much detail or context as timedatectl. Option D is incorrect. Using ls /usr/share/zoneinfo/America/New_York simply lists the file, if it exists, without indicating whether it's the currently active timezone on the system. It does not verify the active system timezone setting.

Question: 5

You are troubleshooting network connectivity issues and need to display the current routing table on a Linux system. Which command should you use?

- A. netstat -r
- B. route -n
- C. ip route show
- D. nmcli route show

Answer: C

Explanation:

ip route show is the modern command to display the routing table on Linux systems. It is part of the iproute2 package, which provides various network management tools. This command gives detailed information about all routes currently known to the kernel. Option A is incorrect. netstat -r can display the routing table, but netstat is considered deprecated in favor of iproute2 tools for modern Linux systems, making ip route show the more current and preferred option. Option B is incorrect. route -n displays the kernel routing table, similar to netstat -r, but it is also part of the older net-tools package and considered less preferred compared to iproute2 tools for routing information. Option D is incorrect. nmcli route show is not a valid nmcli command. While nmcli is used for managing NetworkManager and can show some network settings, displaying the kernel routing table is not one of its functions. The ip route show command is specifically designed for this purpose.

Question: 6

You want to save the output of a command uptime to a file named 'uptime_record.txt' and also display it on the console at the same time. Which command should you use?

- A. uptime > uptime_record.txt
- B. uptime | tee uptime_record.txt
- C. uptime >> uptime_record.txt
- D. uptime &> uptime_record.txt

Answer: B

Explanation:

The tee command reads from standard input and writes to both standard output (the console) and the file named 'uptime_record.txt'. By piping (|) the output of uptime into tee, you achieve the dual goals of saving the output to a file while also displaying it on the console. Option A is incorrect. Using > redirects the output of uptime to 'uptime_record.txt', but it does not also display the output on the console. Option C is incorrect. The >> operator appends the output to 'uptime_record.txt', if it exists, without displaying it on the console. It's used for adding to a file, not for simultaneous display. Option D is incorrect. &> redirects both standard output and standard error to a file, but like >, it does not display the output on the console. This redirection is more about handling error messages along with standard output.

Question: 7

To ensure quick access to a frequently updated data file named 'data_current_month.csv' from another directory, you decide to create a hard link named 'data_link.csv' in that directory. What is the correct command to achieve this?

- A. ln data_current_month.csv /path/to/directory/data_link.csv
- B. ln -s data_current_month.csv /path/to/directory/data_link.csv
- C. cp data_current_month.csv /path/to/directory/data_link.csv
- D. link data_current_month.csv /path/to/directory/data_link.csv

Answer: A

Explanation:

The command ln data_current_month.csv /path/to/directory/data_link.csv correctly creates a hard link named 'data_link.csv' in the specified directory that points to 'data_current_month.csv'. This allows for two or more filenames to refer to the same file on disk. Option B is incorrect. The -s option with ln creates a symbolic link, not a hard link. While symbolic links are useful for many purposes, they do not meet the requirement for a hard link in this scenario. Option C is incorrect. Using cp would make a copy of the file, not create a link. Copies are separate files and would not automatically update when the original file is updated. Option D is incorrect. As mentioned previously, link is not a command typically used directly in shell for creating links; the correct command for creating hard links is ln without any options.

Question: 8

A developer is working on a new feature in a local branch called feature-branch. Before pushing the branch to the remote repository for the first time, the developer wants to ensure their branch includes all recent changes from the main branch. Which Git command should they use to update their local feature-branch with changes from main?

- A. git merge main while on feature-branch
- B. git pull origin main while on main branch
- C. git fetch origin followed by git rebase origin/main while on feature-branch
- D. git checkout main followed by git pull and then git checkout feature-branch and git merge main

Answer: C

Explanation:

This approach ensures that the developer fetches the latest changes from the remote main branch and then rebases their feature-branch onto these updates. This method is preferred in many workflows as it keeps the project history clean and linear, making it easier to understand the sequence of changes.

Option A is incorrect. Using `git merge main` directly merges the latest changes from main into feature-branch, but it can create a merge commit, potentially cluttering the project history if not desired in certain workflows. Option B is incorrect. `git pull origin main` while on the main branch updates the main branch but does not update the feature-branch with the latest changes from main. Option D is incorrect. This method achieves the goal of updating feature-branch with changes from main, but it's more cumbersome and involves unnecessary switching between branches. The rebase approach (Option C) is cleaner and maintains a linear history.

Question: 9

To display the number of lines in 'document.txt' that do not contain the word 'confidential', which command would you use?

- A. `grep -v 'confidential' document.txt | wc -l`
- B. `grep 'confidential' document.txt > wc -l`
- C. `grep -c 'confidential' document.txt`
- D. `cat document.txt | grep -v 'confidential' > wc -l`

Answer: A

Explanation:

This command uses `grep` with the `-v` option to invert the match, selecting lines that do not contain 'confidential'. The output is then piped to `wc -l`, which counts the number of lines. This approach directly answers the requirement to count lines without 'confidential'. Option B is incorrect. This syntax is incorrect because it attempts to redirect the output of `grep` to a command `wc -l` as if it were a file, which is not a valid operation. Option C is incorrect. The `grep -c` option counts the number of lines that match the given pattern, which is the opposite of the requirement to count lines that do not contain 'confidential'. Option D is incorrect. This approach incorrectly attempts to redirect output to `wc -l` as if it were a file and does not properly pipe the output of `grep` to `wc -l` for line counting.

Question: 10

A development team is utilizing GitLab for their version control and collaboration. They have implemented a policy requiring that all new features must undergo a thorough code review before being merged into the main branch. A junior developer has just completed work on a new feature and submitted a merge request (MR). What is the next appropriate action according to GitOps best practices for performing change/code review?

- A. Merge the MR immediately to avoid delays in the feature deployment process

- B. Bypass the code review process for minor features to expedite integration
- C. Assign the MR to at least one senior developer for review and approval before merging
- D. Automatically approve and merge the MR using a bot to streamline the process

Answer: C

Explanation:

Assigning the merge request to a senior developer ensures that the new feature is thoroughly reviewed for quality, security, and alignment with project standards before it is integrated into the main branch. This practice is central to GitOps, emphasizing collaboration, transparency, and quality assurance in the deployment pipeline. Option A is incorrect. Merging the MR immediately without review goes against the GitOps principle of ensuring quality and adherence to standards through peer reviews before changes are deployed. Option B is incorrect. Bypassing the code review process, even for minor features, compromises the integrity of the codebase and neglects the opportunity for learning and improvement, contrary to GitOps practices. Option D is incorrect. Automatically approving and merging MRs without human review contradicts the GitOps emphasis on thorough, collaborative reviews to maintain code quality and security.

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>