Exam : **XK0-006**

Title : CompTIA Linux+ Certification Exam

Vendor : CompTIA

Version : DEMO

NO.1 A systems administrator is decommissioning a service. Which of the following commands should the administrator use to make sure users cannot start the service again?

A. systemctl mask service

B. systemctl kill service

C. systemctl isolate service

D. systemctl disable service

Answer: A

Explanation:

The correct answer is A. systemctl mask service because masking a service ensures that it cannot be started manually or automatically under any circumstances. When a service is masked, its unit file is linked to /dev
/null, effectively making it impossible for systemd to start the service. This is the most appropriate action when permanently decommissioning a service and ensuring that no user or process can restart it.

Option D (systemctl disable service) is a common distractor but is incorrect in this context. Disabling a service only prevents it from starting automatically at boot time; however, users with sufficient permissions can still manually start the service using systemctl start. Therefore, it does not fully meet the requirement of preventing the service from being started again.

Option B (systemctl kill service) is incorrect because it only terminates the currently running instance of the service. It does not prevent the service from being restarted later, either manually or automatically.

Option C (systemctl isolate service) is also incorrect. The isolate command is used to switch the system to a different target (similar to changing runlevels), not to manage the long-term availability of a specific service.

In Linux system administration, particularly within the Linux+ objectives, understanding the difference between stop, disable, and mask is critical. While stopping halts a service temporarily and disabling prevents automatic startup, masking is the strongest control, ensuring the service cannot be activated at all. This makes it the correct and secure choice when decommissioning services in production environments.

NO.2 A Linux administrator just finished setting up passwordless SSH authentication between two nodes. However, upon test validation, the remote host prompts for a password. Given the following logs:

```
-rw-------. 1 root root 588 Apr 3 2022 authorized_keys

avc: denied { read } for pid=xxxx comm="sshd" name="authorized_keys" dev="dm-5" ino=xxxx scontext=system_u:system_r:sshd_t:s0-s0:c0.c1023
tcontext=unconfined_u:object_r:home_root_t:s0 tclass=file
[...]

SELinux status: enabled
SELinuxfs mount: /sys/fs/selinux
SELinux root directory: /etc/selinux
Loaded policy name: targeted
Current mode: enforcing
Mode from config file: enforcing
Policy MLS status: enabled
Policy deny_unknown status: allowed
Max kernel policy version: 31
```

Which of the following is the most likely cause of the issue?

A. The SELinux policy is incorrectly targeting the unconf ined_u context.

B. The administrator forgot to restart the SSHD after creating the authorizedjceys file.

C. The authorized_keys file has the incorrect root permissions assigned.

D. The authorized_keys file does not have the correct security context to match SELinux policy.

Answer: D
Explanation:
This issue is directly related to SELinux enforcement, which is a key topic in the Security domain of CompTIA Linux+ V8. The logs clearly indicate that SSH key-based authentication is failing due to an SELinux access control violation rather than a traditional file permission or SSH configuration problem.
The most important clue is the AVC denial message, which shows that the sshd process is being denied read access to the authorized_keys file. The security context of the file is listed as unconfined_u:object_r:
home_root_t:s0. Under a targeted SELinux policy, SSH is only permitted to read authorized_keys files that are labeled with the correct SELinux type, typically ssh_home_t.
Because SELinux is running in enforcing mode, it actively blocks access that violates policy rules, even if standard UNIX permissions are correct. Although the file permissions (600) are acceptable for an authorized_keys file, SELinux does not rely solely on traditional permissions. The mismatch between the expected SELinux context and the actual context prevents sshd from accessing the file, causing SSH to fall back to password authentication.
Option D correctly identifies the root cause: the authorized_keys file does not have the correct SELinux security context. This is a well-documented Linux+ V8 troubleshooting scenario, commonly resolved by restoring the correct context using commands such as restorecon or by ensuring the file resides in a properly labeled home directory.
The other options are incorrect. Restarting sshd does not fix SELinux labeling issues. The policy itself is functioning as intended, and file ownership alone does not override SELinux access controls. Linux+ V8 documentation emphasizes that SELinux denials must be addressed by correcting file contexts rather than weakening security controls. Therefore, the correct answer is D.

NO.3 An administrator wants to search a file named myFile and look for all occurrences of strings containing at least five characters, where characters two and five are i, but character three is not b. Which of the following commands should the administrator execute to get the intended result?
A. grep .a*^b-.a myFile
B. grep .a., [a] myFile
C. grepa^b*a myFile
D. grep .i[^b].i myFile
Answer: D
Explanation:
Pattern matching using regular expressions is a key troubleshooting and text-processing skill covered in CompTIA Linux+ V8. The grep command, combined with regular expressions, allows administrators to search for complex string patterns within files.
The requirement specifies:
The string must contain at least five characters
Character 2 must be i
Character 3 must not be b
Character 5 must be i
To meet these conditions, the correct regular expression structure is:
any character (position 1)
i # literal i (position 2)
[^b] # any character except b (position 3)

any character (position 4)
i # literal i (position 5)
This results in the expression:
i[^b].i
Option D, grep .i[^b].i myFile, correctly implements this logic. It ensures positional matching and excludes unwanted characters using a negated character class ([^b]), which is explicitly covered in Linux+ V8 regular expression objectives.
The other options contain invalid or malformed regular expressions and do not meet the positional or exclusion requirements. Linux+ V8 emphasizes understanding anchors, character classes, and position-based matching when troubleshooting log files or configuration data.
Therefore, the correct answer is D.

NO.4 A small group of remote users needs permission to run some administrative commands on a Linux server.
Which of the following approaches is the best way to address this request?
A. Set the sticky bit in all binaries requested for use by this group of users.
B. Request a list of utilities they need to use with elevated privileges, and add the appropriate rules.
C. Set the configuration parameter PermitRootLogin on sshd_config to " yes " and share the root password with this group of users.
D. Create a Linux group, add the users, and configure this group on /etc/sudoers to use sudo, thus allowing them to switch to root user.
Answer: D
Explanation:
The correct answer is D. Create a Linux group, add the users, and configure this group on /etc/sudoers to use sudo because it follows the principle of least privilege while providing controlled administrative access. In Linux systems, sudo is the recommended mechanism for delegating elevated privileges without granting full root access.
By creating a dedicated group and adding users to it, administrators can centrally manage permissions.
Configuring this group in the /etc/sudoers file (preferably using visudo) allows fine-grained control over which commands can be executed with elevated privileges. This approach ensures accountability, as sudo logs all executed commands, and enhances security by avoiding password sharing.
Option A is incorrect because the sticky bit does not grant administrative privileges; it is used primarily on directories to restrict file deletion.
Option B is partially correct in principle (least privilege), but it is incomplete because it does not specify implementation. Without integrating those rules into sudoers or another privilege delegation mechanism, it does not provide a complete solution.
Option C is incorrect and highly insecure. Enabling root login over SSH and sharing the root password violates security best practices, eliminates accountability, and significantly increases the risk of compromise.
From a Linux+ security perspective, using sudo with group-based access control is the best practice for privilege delegation. It ensures secure, auditable, and manageable administrative access, aligning with enterprise security standards and minimizing risk exposure.

NO.5 Which of the following is a characteristic of Python 3?

A. It is closed source.

B. It is extensible through modules.

C. It is fully backwards compatible.

D. It is binary compatible with Java.

Answer: B

Explanation:

Python 3 characteristics are part of Linux+ V8 scripting objectives. One of Python's most important features is its modular and extensible architecture.

Option B is correct because Python 3 supports extensibility through modules and packages. Python includes a large standard library and allows developers to extend functionality using third-party modules or custom code.

This makes Python highly adaptable for automation, system management, and DevOps tasks.

The other options are incorrect. Python is open source, not closed source. Python 3 is not fully backwards compatible with Python 2, which is a major distinction emphasized in Linux+ V8. Python is also not binary compatible with Java.

Linux+ V8 documentation highlights Python's extensibility as a key reason it is widely used in Linux automation. Therefore, the correct answer is B.

NO.6 An administrator wants to start an Ubuntu terminal in a container. Which of the following will allow the administrator to complete the task?

A. podman pull ubuntu /bin/bash

B. docker run -it ubuntu

C. containerd /usr/bin/sh ubuntu

D. runc exec -it ubuntu

Answer: B

Explanation:

The correct answer is B. docker run -it ubuntu because it is the standard command used to start an interactive terminal session inside a new Ubuntu container. The docker run command creates and starts a container from a specified image. The -it flags are critical here: -i keeps STDIN open (interactive mode), and -t allocates a pseudo-terminal, allowing the user to interact with the container as if it were a regular shell session.

When executed, Docker will first check if the ubuntu image exists locally. If it does not, Docker will automatically pull it from the default registry (Docker Hub). Once the image is available, the container is launched and provides access to a shell (typically /bin/bash or /bin/sh) inside the Ubuntu environment.

Option A is incorrect because podman pull ubuntu /bin/bash is not valid syntax. The pull command is only used to download images and does not execute a shell.

Option C is incorrect because containerd is a low-level container runtime and is not typically used directly with such syntax to start interactive containers.

Option D is incorrect because runc exec is used to execute commands in an already running container, not to start a new container. Additionally, it requires a container ID rather than an image name.

From a Linux+ perspective, understanding container lifecycle commands is essential for automation and orchestration. The docker run -it command is fundamental for testing, debugging, and interacting with containerized environments, making it a key skill for administrators working with modern

infrastructure.

NO.7 Which of the following does dmesg display?
A. Incorrect login attempts
B. " Session closed " messages
C. Window manager warnings
D. USB device connections
Answer: D
Explanation:
The dmesg (diagnostic message) utility is a critical tool for system management and hardware troubleshooting in Linux+ V8. It is used to display the kernel ring buffer, which contains messages generated by the Linux kernel during the boot process and while the system is running.
The kernel ring buffer primarily records events related to hardware initialization, device driver status, and system resources. When a hardware device is connected to the system-such as a USB drive, a network card, or a keyboard-the kernel detects the event and logs the details. For example, when a USB device is plugged in, dmesg will show the manufacturer, the device ID, and the mount point or device node (e.g., /dev/sdb1) assigned to it.
The other options refer to logs that are typically handled by different services:
* Incorrect login attempts (Option A) and " Session closed " messages (Option B) are authentication and security events. These are usually logged by sshd or pam and stored in /var/log/auth.log or /var/log
/secure.
* Window manager warnings (Option C) are related to the graphical user interface (X11 or Wayland) and are stored in desktop-specific log files or the systemd journal.
According to CompTIA Linux+ documentation, dmesg is the primary tool for verifying hardware detection and diagnosing driver issues. Because it focuses on the kernel-hardware interface, " USB device connections " is the correct answer among the choices provided.

NO.8 Which of the following commands should an administrator use to convert a KVM disk file to a different format?
A. qemu-kvm
B. qemu-nq
C. qemu-io
D. qemu-img
Answer: D
Explanation:
The correct answer is D. qemu-img because it is the standard utility used to create, convert, and manage disk image files used by virtualization platforms such as KVM (Kernel-based Virtual Machine). One of its primary functions is converting disk images between different formats, such as qcow2, raw, vmdk, and vdi.
For example, an administrator can convert a disk image using a command like:
qemu-img convert -f qcow2 -O raw disk.qcow2 disk.raw
This capability is essential when migrating virtual machines between different hypervisors or when optimizing storage formats.
Option A (qemu-kvm) is incorrect because it is used to run virtual machines using KVM acceleration,

not for managing or converting disk images.
Option B (qemu-nq) is incorrect because it is not a valid or commonly used QEMU command.
Option C (qemu-io) is incorrect because it is primarily used for debugging and testing disk I/O operations, not for format conversion.
From a Linux+ system management perspective, virtualization is a key topic, and tools like qemu-img are essential for managing virtual disk storage. Administrators frequently use it to resize images, check integrity, and convert formats when working in cloud, containerized, or virtualized environments. Understanding this command ensures efficient handling of virtual machine storage and supports interoperability across virtualization platforms.

NO.9 A Linux administrator updates the DNS record for the company using:
cat /etc/bind/db.abc.com
The revised partial zone file is as follows:
ns1 IN A 192.168.40.251
ns2 IN A 192.168.40.252
www IN A 192.168.30.30
When the administrator attempts to resolve www.abc.com to its IP address, the domain name still points to its old IP mapping:
nslookup www.abc.com
Server: 192.168.40.251
Address: 192.168.40.251#53
Non-authoritative answer:
Name: www.abc.com
Address: 199.168.20.81
Which of the following should the administrator execute to retrieve the updated IP mapping?

A. systemd-resolve query www.abc.com

B. systemd-resolve status

C. service nslcd reload

D. resolvectl flush-caches

Answer: D

Explanation:
This scenario represents a classic DNS troubleshooting situation covered in the Troubleshooting domain of the CompTIA Linux+ V8 objectives. Although the DNS zone file has been updated correctly on the BIND server, the system continues to resolve the domain name to an outdated IP address. This behavior strongly indicates DNS caching rather than a configuration error in the zone file itself.
Modern Linux systems that use systemd-resolved cache DNS responses locally to improve performance and reduce external queries. Even after a DNS record is updated on the authoritative server, cached results may persist until the cache expires or is manually cleared. The nslookup output showing a non-authoritative answer further confirms that the response is being served from a cache rather than directly from the updated zone data.
The correct solution is to flush the local DNS cache so the system can retrieve the updated record from the DNS server. The command resolvectl flush-caches clears all cached DNS entries maintained by systemd- resolved, forcing fresh queries to authoritative name servers. This aligns directly with Linux+ V8 documentation for resolving name resolution inconsistencies caused by stale cache entries.
The other options are incorrect for the following reasons. systemd-resolve query www.abc.com performs a DNS lookup but does not clear cached entries. systemd-resolve status only displays

resolver configuration and statistics. service nslcd reload reloads the Name Service LDAP daemon and is unrelated to DNS resolution or caching.
Linux+ V8 emphasizes identifying whether issues originate from services, configuration, or cached data. In this case, flushing the DNS cache is the correct and least disruptive corrective action.
Therefore, the correct answer is D. resolvectl flush-caches.

NO.10 The development team asks a Linux administrator to help diagnose a connectivity issue that is occurring with a newly developed software. The Linux administrator reviews the following output:
$ wget -vvv https://api.newapp.comptia.org/v2/health
Resolving proxy.comptia.org (proxy.comptia.org)... connected.
ERROR: The certificate of ' api.newapp.comptia.org ' is not trusted.
ERROR: The certificate of ' api.newapp.comptia.org ' does not have a known issuer.
Which of the following actions is the best way to resolve the issue?

A. Verify the remote certificate is trustworthy, and add it to the local trusted certificates repository.

B. Replace the remote certificate with a new self-signed CA and application certificate.

C. Replace the expired CA certificate to permanently resolve the problem.

D. Add option --no-check-certificate to wget command to permanently resolve the problem.

Answer: A

Explanation:
The correct answer is A. Verify the remote certificate is trustworthy, and add it to the local trusted certificates repository because the error clearly indicates that the system does not recognize the certificate authority (CA) that issued the server's SSL/TLS certificate. This is a trust issue, not necessarily a problem with the certificate itself.
When wget reports that the certificate "is not trusted" and "does not have a known issuer," it typically means the CA certificate is missing from the local trust store. The proper and secure resolution is to first validate that the remote certificate is legitimate (for example, confirming it with the issuing authority or organization).
Once verified, the administrator should add the CA certificate to the system's trusted certificate store (e.g., /etc
/pki/ca-trust/ or /usr/local/share/ca-certificates/ depending on the distribution) and update the trust database.
Option B is incorrect because using a self-signed certificate introduces additional trust issues and is not appropriate for production environments unless properly managed.
Option C is incorrect because the error message does not indicate that the certificate is expired, only that the issuer is unknown.
Option D is incorrect because using --no-check-certificate bypasses SSL verification entirely, which creates a significant security vulnerability and violates best practices.
From a Linux+ security perspective, maintaining proper certificate validation is essential for secure communications. Administrators must ensure that trusted CAs are properly configured rather than bypassing verification mechanisms.

NO.11 A Linux administrator wants to use AI to deploy infrastructure as code. Which of the following is a best practice regarding the use of AI for this task?

A. Using copy and paste when possible

B. Generating monolithic code

C. Linting generated code
D. Merging CI/CD pipelines
Answer: C
Explanation:
The use of AI-assisted tools for Infrastructure as Code (IaC) is increasingly common in modern DevOps practices, which are covered in the Automation and Orchestration domain of CompTIA Linux+ V8. While AI can accelerate code generation, Linux+ emphasizes that generated code must still follow best practices for quality, security, and maintainability.
Option C, linting generated code, is the correct answer. Linting involves analyzing code for syntax errors, style violations, logical issues, and potential security risks before deployment. When AI tools generate IaC artifacts such as Terraform files, Ansible playbooks, or shell scripts, linting ensures that the output adheres to organizational standards and does not introduce misconfigurations or vulnerabilities.
Linux+ V8 documentation stresses that automation does not eliminate the need for validation. Administrators remain responsible for verifying correctness and compliance. Linting tools such as ansible-lint, terraform fmt, and shell linters help detect issues early in the pipeline and prevent faulty infrastructure deployments.
The other options are poor practices. Copying and pasting code increases the risk of errors and inconsistency.
Generating monolithic code contradicts modular IaC principles. Merging CI/CD pipelines is unrelated to validating AI-generated infrastructure code.
Therefore, the correct answer is C. Linting generated code.

NO.12 A Linux administrator wants to make the enable_auth variable set to 1 and available to the environment of subsequently executed commands. Which of the following should the administrator use for this task?
A. let ENABLE_AUTH=1
B. ENABLE_AUTH=1
C. ENABLE_AUTH=$(echo $ENABLE_AUTH)
D. export ENABLE_AUTH=1
Answer: D
Explanation:
Environment variables in Linux can exist either locally within a shell or be exported to child processes. CompTIA Linux+ V8 emphasizes the distinction between shell variables and environment variables, as this affects how applications inherit configuration values.
Option D, export ENABLE_AUTH=1, is the correct choice because it both assigns the variable and marks it for export to the environment. Once exported, the variable becomes available to all subsequently executed commands and child processes spawned from the current shell. This behavior is required when applications or scripts rely on environment variables for configuration.
Option B, ENABLE_AUTH=1, only sets a shell-local variable. While it is accessible within the current shell session, it is not inherited by child processes unless explicitly exported. Option A, let ENABLE_AUTH=1, performs arithmetic evaluation and does not export the variable. Option C incorrectly assigns the output of a command substitution and does not set the desired value.
Linux+ V8 documentation highlights export as the correct mechanism for making variables available system- wide within a user session. Therefore, the correct answer is D.

NO.13 An administrator set up a new user account called " test " . However, the user is unable to change their password. Given the following output:

```
[test@localhost ~]$ passwd
Changing password for user test.
Current password:
passwd: Authentication token manipulation error

[test@localhost ~]$ ls -l /bin/passwd
-rwxr-xr-x. 1 root root 33424 Feb 7 2022 /bin/passwd

[test@localhost ~]$ chage -l test
Last password change : Oct 19, 2023
Password expires : never
Password inactive : never
Account expires : never
Minimum number of days between password change : 0
Maximum number of days between password change : 99999
Number of days of warning before password expires : 7

[root@localhost bin]# passwd test
Changing password for user test.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
```

Which of the following is the most likely cause of this issue?

A. The SUID bit is missing on the /bin/passwd file.

B. The password provided by the user " test " does not meet complexity requirements.

C. The user " test " already changed the password today.

D. The password has been disabled for user " test " .

Answer: A

Explanation:

For normal users to change their own password, /bin/passwd must have the SUID bit set (permissions should be -rwsr-xr-x). The SUID bit allows users to run the program with the permissions of the file owner (root), which is required to update /etc/shadow. The provided output shows /bin/passwd does not have the SUID bit (no ' s ' in the owner's execute field). As a result, user " test " receives an " Authentication token manipulation error " . The password can be changed as root, which confirms it ' s a permissions/SUID issue.

Other options:

B). If the password didn't meet requirements, a different error would appear.

C). There is no minimum day limit preventing password change (see chage -l output).

D). The account and password are active (not disabled).

Reference:

CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 6: " User and Group Management " , Section: " Managing User Passwords and Policies " CompTIA Linux+ XK0-006 Objectives, Domain 1.0: System Management

NO.14 An administrator updates the network configuration on a server but wants to ensure the change will not cause an outage if something goes wrong. Which of the following commands allows the administrator to accomplish this goal?

A. netplan try

B. netplan rebind

C. netplan ip

D. netplan apply

Answer: A

Explanation:

Network configuration changes can cause immediate loss of connectivity if applied incorrectly. Linux+ V8 emphasizes safe configuration practices, particularly when managing remote systems.

The netplan try command applies network configuration changes temporarily and prompts the administrator to confirm them within a timeout period. If the administrator does not confirm, Netplan automatically rolls back to the previous working configuration. This prevents accidental outages caused by misconfigured network settings.

The netplan apply command makes changes permanent immediately and does not provide rollback protection.

The other options are not valid Netplan commands.

Linux+ V8 documentation explicitly references netplan try as a safe testing mechanism. Therefore, the correct answer is A.

NO.15 A systems administrator is helping to secure a new web application. During the tests, the administrator obtains the following output to validate the application:

SSL connection using TLSv1.2 / ECDHE-RSA-AES256-GCM-SHA384

ALPN, server accepted to use http/1.1

Server certificate:

subject: CN=*.newapp.comptia.org

start date: Jan 17 00:00:00 2024 GMT

expire date: Feb 16 23:59:59 2034 GMT

issuer: C=US; O=Comptia; OU=IT Security; CN=ca1.comptia.org

SSL certificate verify result: unable to get local issuer certificate (20), continuing anyway.

Which of the following explains the validation results?

A. The certificate was not signed by a trusted authority, which was forcefully ignored during the tests.

B. The certificate was only issued for one month, is already expired, and needs to be replaced.

C. The certificate is a wildcard certificate that is highly insecure and should never be used.

D. The certificate uses an outdated algorithm and should be replaced with a more secure one.

Answer: A

Explanation:

The correct answer is A. The certificate was not signed by a trusted authority, which was forcefully ignored during the tests. The key indicator in the output is the message: "SSL certificate verify result:

unable to get local issuer certificate (20), continuing anyway." This error means that the client system cannot validate the certificate chain because it does not recognize or trust the certificate authority (CA) that issued the server certificate.
In TLS/SSL validation, the client must verify that the server's certificate chains up to a trusted root CA present in the local trust store. If the issuer (in this case, ca1.comptia.org) is not included in the client's trusted certificate authorities, the verification fails. However, the phrase "continuing anyway" indicates that the validation failure was bypassed-likely due to a testing flag such as -k or --insecure in curl-allowing the connection despite the trust issue.
Option B is incorrect because the certificate validity dates show it is valid from 2024 to 2034, not expired.
Option C is incorrect because wildcard certificates (e.g., *.newapp.comptia.org) are commonly used and not inherently insecure when properly managed.
Option D is incorrect because the cipher suite shown (ECDHE-RSA-AES256-GCM-SHA384) is considered strong and modern, not outdated.
From a Linux+ security perspective, proper certificate validation is critical. Administrators should ensure that the issuing CA is trusted by installing the correct CA certificates rather than bypassing validation, which exposes systems to man-in-the-middle attacks and other security risks.