

# CompTIA

## Exam Questions XK0-006

CompTIA Linux+ Exam



### NEW QUESTION 1

A Linux administrator receives reports about MySQL service availability issues. The administrator observes the following information:

- uptime -p shows the system has been up for only 2 minutes
- journalctl shows messages indicating:mysqld invoked oom-killermysqld cpuset=/ mems\_allowed=0 Which of the following explains why the server was offline?

- A. The process exhausted server memory.
- B. The process was intentionally terminated by a privileged user.
- C. The process crashed because of a filesystem error.
- D. A network outage caused a service availability issue.

**Answer:** A

#### Explanation:

is A. The process exhausted server memory.

### NEW QUESTION 2

A systems administrator needs to integrate a new storage array into the company's existing storage pool. The administrator wants to ensure that the server is able to detect the new storage array. Which of the following commands should the administrator use to ensure that the new storage array is presented to the systems?

- A. lsscsi
- B. lsusb
- C. lspci
- D. lshw

**Answer:** A

#### Explanation:

Comprehensive and Detailed Explanation: From Exact Extract:

The lsscsi command is used to list information about SCSI devices (including storage arrays) that are attached to the system. This is critical when integrating a new storage array because it allows the administrator to verify that the operating system detects the new device at the SCSI layer, which is the underlying interface for most enterprise storage solutions. lsscsi outputs a list of recognized SCSI devices, their device nodes, and associated information.

Other options:

- B. lsusb: Lists USB devices, not storage arrays on SCSI/SATA/SAS.
- C. lspci: Displays information on IPC (inter-process communication) facilities, unrelated to hardware detection.
- D. lshw: Lists hardware details and can show storage, but lsscsi is specifically designed for SCSI device detection and is the most direct method for this task.

Reference:

CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 7: "Managing Storage", Section: "Identifying and Accessing Storage Devices"

CompTIA Linux+ XK0-006 Objectives: Domain 4.0 – Storage and Filesystems

=====

### NEW QUESTION 3

A systems administrator is reconfiguring existing user accounts in a Linux system. Which of the following commands should the administrator use to include "myuser" in the finance group?

- A. groupadd finance myuser
- B. groupmod finance myuser
- C. useradd -g finance myuser
- D. usermod -aG finance myuser

**Answer:** D

#### Explanation:

Comprehensive and Detailed Explanation: From Exact Extract:

To add an existing user (myuser) to an existing group (finance) without removing them from other groups, the correct command is usermod -aG finance myuser.

The -aG option appends the user to the supplementary group

(s) specified.

Other options:

- A. groupadd is for creating new groups, not adding users to groups.
- B. groupmod is for modifying group properties, not user membership.
- C. useradd creates new users; not applicable to existing users.

Reference:

CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 6: "User and Group Management", Section: "Modifying Group Membership"

CompTIA Linux+ XK0-006 Objectives, Domain 1.0: System Management

=====

### NEW QUESTION 4

A Linux administrator needs to create and then connect to the app-01-image container. Which of the following commands accomplishes this task?

- A. docker run -it app-01-image

- B. docker start -td app-01-image
- C. docker build -ic app-01-image
- D. docker exec -dc app-01-image

**Answer:** A

**Explanation:**

Comprehensive and Detailed 250 to 350 words of Explanation: From Linux+ V8 documents:

Container lifecycle management is a core topic within the Automation, Orchestration, and Scripting domain of CompTIA Linux+ V8. Administrators must understand the difference between creating containers, starting containers, and executing commands within running containers.

The correct command is docker run -it app-01-image. The docker run command performs three actions at once: it creates a new container from the specified image, starts the container, and optionally attaches the administrator's terminal to it. The -i option keeps standard input open, while the -t option allocates a pseudo-terminal (TTY). Together, these options allow the administrator to interactively connect to the container immediately after it is created.

The other options are incorrect for the following reasons. docker start is used only to start an existing stopped container and does not create a new container from an image. Additionally, -t and -d are not valid options for attaching an interactive terminal during container startup. docker build is used to build a Docker image from a Dockerfile and cannot be used to create or connect to a container. docker exec is used to run commands inside an already running container and therefore cannot be used to create a container.

Linux+ V8 documentation emphasizes that docker run is the primary command used when administrators want to instantiate containers from images and interact with them. This command is commonly used during testing, development, and troubleshooting workflows.

**NEW QUESTION 5**

A systems administrator is preparing a Linux system for application setup. The administrator needs to create an environment variable with a persistent value in one of the user accounts. Which of the following commands should the administrator use for this task?

- A. export "VAR=SomeValue" >> ~/.ssh/profile
- B. export VAR=value
- C. VAR=value
- D. echo "export VAR=value" >> ~/.bashrc

**Answer:** D

**Explanation:**

Environment variables are widely used in Linux systems to configure application behavior, and Linux+ V8 emphasizes the distinction between temporary and persistent variables. A variable is persistent only if it is defined in a shell initialization file.

The correct approach is echo "export VAR=value" >> ~/.bashrc. This command appends the variable definition to the user's .bashrc file, ensuring the variable is set automatically every time the user starts a new shell session. This makes the variable persistent for that specific user.

Options B and C only define variables in the current shell session and are lost when the session ends. Option A incorrectly targets the SSH configuration directory and is not appropriate for defining shell environment variables.

Linux+ V8 documentation highlights .bashrc, .bash\_profile, and /etc/profile as correct locations for persistent environment variables. Therefore, the correct answer is D.

**NEW QUESTION 6**

Which of the following most accurately describes a webhook?

- A. An authentication method for web-server communication
- B. An SNMP-based API for network device monitoring
- C. A means to transmit sensitive information between systems
- D. An HTTP-based callback function

**Answer:** D

**Explanation:**

Webhooks are commonly used in automation and DevOps workflows, which are emphasized in the Linux+ V8 objectives. A webhook is best described as an HTTP-based callback mechanism that allows one system to notify another when a specific event occurs.

Option D correctly defines a webhook. Instead of polling an API at regular intervals, a webhook allows an application to automatically send an HTTP request—typically a POST—to a predefined URL when an event happens. This makes webhooks efficient, event-driven, and well-suited for automation pipelines, CI/CD systems, and monitoring integrations.

The other options are incorrect. Option A confuses webhooks with authentication mechanisms. Option B incorrectly associates webhooks with SNMP, which is a separate protocol. Option C is misleading because webhooks are not inherently designed for transmitting sensitive data and require additional security measures such as TLS and authentication.

Linux+ V8 documentation highlights webhooks as a key integration method in automated environments, enabling systems to react in real time to changes or triggers.

Therefore, the correct answer is D.

**NEW QUESTION 7**

In the echo "profile-\$num-\$name" line of a shell script, the variable \$num seems to not be expanding during execution. Which of the following notations ensures the value is expanded?

- A. echo "profile-\$(num)-\$name"
- B. echo 'profile-\$num-\$name'
- C. echo "profile-'\$num'-\$name"
- D. echo "profile-\${num}-\$name"

**Answer:** D

**Explanation:**

Shell variable expansion is a fundamental scripting concept included in Linux+ V8 objectives. In Bash and similar shells, variables are expanded only when they are interpreted within double quotes or unquoted contexts, and sometimes explicit syntax is required to avoid ambiguity.

The correct notation is \${num}, as shown in option D. Using curly braces around the variable name ensures the shell correctly identifies the variable boundary, especially when it is adjacent to other characters. This guarantees proper expansion of the variable's value.

The other options are incorrect. Single quotes prevent variable expansion entirely. The `$(...)` syntax is used for command substitution, not variable expansion. Quoting the variable name itself also prevents expansion. Linux+ V8 documentation emphasizes `${VAR}` notation as a best practice in shell scripting for clarity and correctness. Therefore, the correct answer is D.

**NEW QUESTION 8**

Which of the following best describes a use case for playbooks in a Linux system?

- A. To provide a set of tasks and configurations to deploy an application
- B. To provide the instructions for implementing version control on a repository
- C. To provide the security information required for a container
- D. To provide the storage volume information required for a pod

**Answer:** A

**Explanation:**

In the context of Linux automation and orchestration, playbooks are most commonly associated with configuration management tools such as Ansible, which is explicitly referenced in the CompTIA Linux+ V8 objectives. Playbooks are written in YAML and are designed to define a series of tasks, configurations, and desired system states that should be applied to one or more Linux systems in a repeatable and automated manner.

A primary use case for playbooks is application deployment and system configuration automation. Playbooks allow administrators to specify tasks such as installing packages, configuring services, managing users, setting permissions, deploying application files, and starting or enabling services. This aligns directly with option A, which accurately describes playbooks as a method to provide a set of tasks and configurations required to deploy an application consistently across environments.

The remaining options are not accurate representations of playbook functionality. Option B refers to version control implementation, which is handled by tools like Git and is not the purpose of playbooks themselves, although playbooks may be stored in version control systems. Option C describes container security information, which is typically managed through container runtime configurations, secrets, or security policies rather than playbooks. Option D refers to storage volume information for a pod, which is specific to Kubernetes manifests and not a general Linux playbook use case.

According to Linux+ V8 documentation, automation tools and playbooks help reduce human error, improve consistency, and support Infrastructure as Code (IaC) practices. Playbooks are a key mechanism for orchestrating multi-step operations across multiple systems, making them essential for modern Linux system administration.

Therefore, the correct answer is A, as it best describes the practical and documented use case for playbooks in a Linux system.

**NEW QUESTION 9**

An administrator is investigating the reason a Linux workstation is not resolving the website `http://www.comptia.org`. The administrator executes some commands and receives the following output:

```
$ dig @8.8.8.8 www.comptia.org +short
104.18.16.29

$ nslookup -querytype=A www.comptia.org
...
Name: www.comptia.org
Address: 104.18.16.29

$ nslookup -querytype=AAAA www.comptia.org
...
*** Can't find www.comptia.org: No answer

$ ping -4 www.comptia.org
PING www.comptia.org (104.18.99.101)
From somehost (192.168.1.192) icmp_seq=3 Destination Host Unreachable
...

$ cat /etc/hosts
127.0.0.1 localhost localhost.localdomain
104.18.99.101 www.comptia.org
```

Which of the following is the most likely cause?

- A. The static entry needs to be removed from `/etc/hosts`.
- B. The remote website does not support IPv6, and the workstation requires it.
- C. The firewall needs to be modified to allow outbound HTTP and HTTPS.
- D. The nameserver in `/etc/resolv.conf` needs to be updated to 8.8.8.8

**Answer:** A

**Explanation:**

When troubleshooting name resolution issues in Linux, `/etc/hosts` entries take precedence over DNS lookups. The workstation's `/etc/hosts` file contains the line: `CopyEdit 104.18.99.101 www.comptia.org`

This means any attempt to access `www.comptia.org` will resolve to `104.18.99.101`, regardless of the real DNS response. However, both `dig` and `nslookup` show the

correct IP as 104.18.16.29. Because the local /etc/hosts entry overrides DNS, and the hardcoded IP is either incorrect or unreachable, all network traffic to www.comptia.org will fail or not reach the intended destination, resulting in the observed connectivity issue (Destination Host Unreachable).

Other options:

- \* B. The lack of IPv6 support is irrelevant since the host is using IPv4 and the DNS queries for IPv4 (A record) are successful.
- \* C. The firewall would block all HTTP/HTTPS connections, but the error shown is a host unreachable, not a port-specific issue.
- \* D. The nameserver is working; both dig and nslookup queries succeed and return the correct A record.

[Reference:, CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 8: "Networking Fundamentals", Section: "Troubleshooting Name Resolution", CompTIA Linux+ XK0-006 Objectives, Domain 2.0: Networking, ]

### NEW QUESTION 10

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.

### NEW QUESTION 10

An administrator must secure an account for a user who is going on extended leave. Which of the following steps should the administrator take? (Choose two)

- A. Set the user's files to immutable.
- B. Instruct the user to log in once per week.
- C. Delete the user's /home folder.
- D. Run the command `passwd -l user`.
- E. Change the date on the /home folder to that of the expected return date.
- F. Change the user's shell to /sbin/nologin.

**Answer:** DF

#### Explanation:

Comprehensive and Detailed 250 to 350 words of Explanation From Linux+ V8 documents:

Securing dormant or temporarily unused user accounts is a best practice emphasized in the Security domain of CompTIA Linux+ V8. When a user goes on extended leave, the goal is to prevent unauthorized access while preserving the user's data and account for future use.

The most effective approach is to disable authentication and interactive login access without deleting the account. Option D, running `passwd -l user`, locks the user's password by prepending an invalid character to the encrypted password in /etc/shadow. This prevents password-based authentication while retaining the account, files, and ownership information. Linux+ V8 documentation highlights password locking as a standard method for temporarily disabling accounts. Option F, changing the user's shell to /sbin/nologin, further strengthens account security by preventing interactive shell access entirely. Even if another authentication mechanism were attempted, the user would be denied a login shell. This is a common defense-in-depth measure and is explicitly referenced in Linux+ V8 objectives for access control and account hardening.

The other options are incorrect or inappropriate. Option A (immutable files) does not prevent account access and may interfere with system operations.

Option B defeats the purpose of securing an inactive account. Option C deletes user data, which is unnecessary and risky. Option E has no security effect, as filesystem timestamps do not control access.

Linux+ V8 stresses that secure account management should be reversible, auditable, and minimally disruptive. Locking the password and disabling the login shell meet these criteria and are commonly used together in enterprise environments.

### NEW QUESTION 11

A Linux administrator tries to install Ansible in a Linux environment. One of the steps is to change the owner and the group of the directory /opt/Ansible and its contents. Which of the following commands will accomplish this task?

- A. `groupmod -g Ansible -n /opt/Ansible`
- B. `chown -R Ansible:Ansible /opt/Ansible`
- C. `usermod -aG Ansible /opt/Ansible`
- D. `chmod -c /opt/Ansible`

**Answer:** B

#### Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

The `chown` command is used to change the owner and group of files and directories. The `-R` (recursive) flag ensures that all contents within the directory are also updated. The correct syntax is `chown -R owner:group directory`. So, `chown -R Ansible:Ansible /opt/Ansible` will change the owner and group for /opt/Ansible and everything inside it to "Ansible".

Other options:



- \* A. `groupmod` is used to modify group properties, not ownership of directories or files.
- \* C. `usermod` is for modifying user properties or group memberships.
- \* D. `chmod` changes permissions, not owner/group.

[Reference:, CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 6: "User and Group Management", Section: "Managing File Ownership and Permissions", CompTIA Linux+ XK0-006 Objectives, Domain 1.0: System Management, ]

### NEW QUESTION 16

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

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



- A. Mastered
- B. Not Mastered

**Answer: A**

#### Explanation:

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

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

```
parted -s /dev/sdc mklabel gpt
```

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

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

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

```
mkfs.ext4 /dev/sdc1
```

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

#### NEW QUESTION 20

A user states that an NFS share is reporting random disconnections. The systems administrator obtains the following information

```
#df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/mapper/fedora-root
15G  15G  204K  100% /
devtmpfs        4.0M  0    4.0M  0%   /dev
tmpfs           2.0G  0    2.0G  0%   /dev/shm
tmpfs           783M  816K  782M  1%   /run
tmpfs           2.0G  0    2.0G  0%   /tmp
/dev/vda2       960M  481M  480M  51%  /boot
10.0.0.1:/nfsdata 4T    3.8T 200G  95%  /share

$ ip -s link show
2: enp1s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP mode DEFAULT group default qlen
link/ether 52:5a:00:f7:27:23 brd ff:ff:ff:ff:ff:ff
RX:  bytes      packets    errors    dropped    missed    mcast
    108487310  149198    9584      40721     0         0
TX:  bytes      packets    errors    dropped    carrier    collsns
    3015941    33656    12780     7854     0         0
```

Which of the following best explains the symptoms that are being reported?

- A. The mount point is incorrect for the NFS share.
- B. The IP address of the NFS share is incorrect.
- C. The filesystem is nearly full and is reporting errors.
- D. The interface is reporting a high number of errors and dropped packets.

**Answer: D**

**Explanation:**

This issue is best analyzed using a layered troubleshooting approach, as recommended in the Troubleshooting domain of CompTIA Linux+ V8. The reported symptom is intermittent or random disconnections from an NFS share, which commonly indicates a network reliability issue rather than a configuration or filesystem problem.

The most critical evidence comes from the output of `ip -s link show`. The network interface `enp1s0` is reporting significant numbers of errors and dropped packets on both the receive (RX) and transmit (TX) paths. High packet loss at the network interface level directly affects protocols like NFS, which rely on stable, continuous TCP/IP communication. When packets are dropped or corrupted, NFS clients may experience timeouts, retransmissions, and apparent disconnections. Although the `df -h` output shows that the NFS filesystem is 95% full, this alone does not typically cause random disconnections. A nearly full filesystem may lead to write failures or performance degradation, but it does not explain intermittent connectivity loss. Linux+ V8 documentation notes that filesystem capacity issues usually present as I/O errors, not transport-layer disconnects.

Options A and B can also be ruled out. If the mount point or IP address were incorrect, the NFS share would fail consistently rather than intermittently. The fact that the share is mounted and accessible confirms that the mount configuration and IP addressing are correct.

Linux+ V8 emphasizes that NFS performance and reliability are highly sensitive to network quality. Packet errors, drops, faulty NICs, cabling issues, duplex mismatches, or driver problems commonly result in unstable NFS behavior.

Therefore, the best Explanation for the reported random disconnections is D. The interface is reporting a high number of errors and dropped packets.

**NEW QUESTION 25**

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`.

**NEW QUESTION 27**

A systems administrator is writing a script to analyze the number of files in the directory `/opt/application`

`/home/`. Which of the following commands should the administrator use in conjunction with `ls -l |` to count the files?

- A. `less`
- B. `tail -f`
- C. `tr -c`
- D. `wc -l`

**Answer:** D

**Explanation:**

Explanation

Comprehensive and Detailed Explanation From Exact Extract:

`wc -l` counts the number of lines of input provided to it, which is commonly used to count the number of files when used with `ls -l` (excluding the header line). For example, `ls -l /opt/application/home/ | wc -l` gives the total count of lines, which corresponds to the number of files and directories (including the total line at the top). Other options:

\* A. `less` is a pager utility.

\* B. `tail -f` shows the end of a file in real time.

\* C. `tr -c` translates or deletes characters, not for counting lines.

Reference:

CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 4: "Working with the Command Line", Section: "Text Processing Commands"

CompTIA Linux+ XK0-006 Objectives, Domain 1.0: System Management

## NEW QUESTION 29

A junior system administrator removed an LVM volume by mistake.

### INSTRUCTIONS

Part 1

Review the output and select the appropriate command to begin the recovery process.

Part 2

Review the output and select the appropriate command to continue the recovery process.

Part 3

Review the output and select the appropriate command to complete the recovery process and access the underlying data.

Part 1
Part 2
Part 3

### > Commands

```
[root@comptiasim ~]# df -h
[root@comptiasim ~]# ls -l /dev | grep -v tty
[root@comptiasim ~]# ls -l /etc/lvm/archive
[root@comptiasim ~]# pvdisplay
[root@comptiasim ~]# pvs
[root@comptiasim ~]# vgcfgrestore --list vg01
[root@comptiasim ~]# vgdisplay
[root@comptiasim ~]# vgs
```

```
[root@comptiasim ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        1.9G   0 1.9G   0% /dev
tmpfs           1.9G   0 1.9G   0% /dev/shm
tmpfs           1.9G  17M 1.9G   1% /run
tmpfs           1.9G   0 1.9G   0% /sys/fs/cgroup
/dev/xvda1      8.0G  1.1G 7.0G  13% /
tmpfs           379M   0 379M   0% /run/user/1000
```

Select the appropriate command to begin the recovery process.

```
[root@comptiasim ~]#
```

Select command

```
lvchange -a y /dev/vg01/lv01
lvconvert --type mirror lv01
pvscan
vgcfgrestore vg01 -f /etc/lvm/archive/vg01_00002-966141411.vg
vgcfgrestore vg01 -f /etc/lvm/backup/vg01
lvchange -a n /dev/vg01/lv01
vgcfgrestore vg01 -i -M /etc/lvm/archive/vg01_00001-810050352.vg
```

Select command

Part 2
Part 2
Part 2

### > Commands

```
[root@comptiasim ~]# blkid
[root@comptiasim ~]# dmesg | tail -20
[root@comptiasim ~]# bbbik
[root@comptiasim ~]# ls /
[root@comptiasim ~]# lvdisplay
[root@comptiasim ~]# lvs
[root@comptiasim ~]# lyecan
[root@comptiasim ~]# pvs
[root@comptiasim ~]# vgs
```

```
[root@comptiasim ~]# blkid
/dev/xvda1: UUID="388a99ed-9486-4a46-aeb6-06eaf6c47675" TYPE="xfs"
/dev/xvdf: UUID="1uyvyk-Ffd0-8rvF-cYba-15ZC-EHRZ-JM3Uhm" TYPE="LVM2 member"
```

Select the appropriate command to continue the recovery process.

```
[root@comptiasim ~]#
```

Select command

```
pvchange -x y /dev/xvdf
lvextend -L v54 vg01/lv01 /dev/xvdf
lvchange -x y /dev/vg01/lv01
mount /dev/vg01/lv01/ /important data
lvchange -a n /dev/vg01/lv01
```

Select command

Part 1

Part 2

Part 3

## > Commands

```
[root@comptiasim ~]# blkid
```

```
[root@comptiasim ~]# cat /etc/fstab
```

```
[root@comptiasim ~]# ls -l /dev/mapper/
```

```
[root@comptiasim ~]# ls -l /
```

```
[root@comptiasim ~]# lsblk
```

```
[root@comptiasim ~]# lvdisplay
```

```
[root@comptiasim ~]# lvscan
```

```
[root@comptiasim ~]# tail -f /var/log/messages
```

```
[root@comptiasim ~]# xfs_repair -n /dev/vg01/lv01
```

```
[root@comptiasim ~]# blkid
```

```
/dev/xvda1:          UUID="388a99ed-9496-4a46-aeb6-06eaf6c47675"
```

```
TYPE="xfs"
```

```
/dev/mapper/vg01-lv01: UUID="c63883e9-ceca-45f4-9ad9-f8d8c1814e7e"
```

```
TYPE="xfs"
```

```
/dev/xvdf:          UUID="1uyvyk-Ffd0-RryF-cYba-152C-EHRZ-UN3UHm"
```

```
TYPE="LVM2_member"
```

Select command

xfs\_repair /dev/vg01/lv01

lvscan -a

mount -a

mount /important\_data /dev/vg01/lv01

xfs\_mdrestore /dev/vg01 /important\_data

```
[root@comptiasim ~]#
```

Select command

Part 1
Part 2
Part 3

```

[root@comptia ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/devtmpfs       1.0G     0 1.0G   0% /dev
tmpfs           1.0G     0 1.0G   0% /tmp
tmpfs           1.0G  116K 1.0G   1% /run
tmpfs           1.0G     0 1.0G   0% /tmp/udevlog
/dev/sda1       1.0G  1.0G  1.0G  100% /
tmpfs           1.0G     0 1.0G   0% /dev/shm

```

```

[root@comptia ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/devtmpfs       1.0G     0 1.0G   0% /dev
tmpfs           1.0G     0 1.0G   0% /tmp
tmpfs           1.0G  116K 1.0G   1% /run
tmpfs           1.0G     0 1.0G   0% /tmp/udevlog
/dev/sda1       1.0G  1.0G  1.0G  100% /
tmpfs           1.0G     0 1.0G   0% /dev/shm

```

Select the appropriate command to continue the recovery process.

```

[root@comptia ~]#

```

Select command

```

mount -t ext4 /dev/sda1
mount -t ext4 /dev/sda1 /mnt
mount -t ext4 /dev/sda1 /mnt
mount -t ext4 /dev/sda1 /mnt
mount -t ext4 /dev/sda1 /mnt

```

visit - <https://www.surepassexam.com>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Part 1 – Begin the recovery process Answer

```
vgcfgrestore vg01 -f /etc/lvm/archive/vg01_00001-810050352.vg
```

Part 2 – Continue the recovery process Answer

```
lvchange -ay /dev/vg01/lv01
```

Part 3 – Complete recovery and access data Answer

```
mount /dev/vg01/lv01 /important_data
```

This performance-based question tests LVM recovery, a critical System Management skill in CompTIA Linux+ V8. The scenario indicates that a logical volume was removed, but the underlying physical volume and volume group metadata still exist.

# Part 1: Restoring Volume Group Metadata

The first screenshot shows that:

- \* Physical volumes (pvdisk, pvs) still exist
- \* The logical volume is missing
- \* /etc/lvm/archive/ contains archived VG metadata

Linux automatically stores backups of LVM metadata in /etc/lvm/archive whenever changes are made. The correct first step is to restore the volume group metadata using:

```
vgcfgrestore vg01 -f /etc/lvm/archive/vg01_00001-810050352.vg
```

This restores the logical volume definitions but does not activate them yet.

This is the only correct starting point in Linux+ V8 recovery workflows.

# Part 2: Activating the Logical Volume

After metadata restoration:

- \* The LV exists but is inactive
- \* blkid shows the LV as TYPE="LVM2\_member"

The logical volume must be activated before it can be mounted: 

```
lvchange -ay /dev/vg01/lv01
```

This makes the LV available under /dev/vg01/lv01. Linux+ explicitly requires LV activation after recovery.

# Part 3: Accessing the Data

The final output shows:

- \* The filesystem type is xfs
- \* The logical volume is now visible

Since there is no indication of filesystem corruption, no repair is required.

The correct final step is to mount the filesystem: 

```
mount /dev/vg01/lv01 /important_data
```

This restores full access to the underlying data.

**NEW QUESTION 32**

A Linux user frequently tests shell scripts located in the /home/user/scripts directory. Which of the following commands allows the user to run the program by invoking only the script name?

- A. `export SHELL=$SHELL=/home/user/scripts`
- B. `export TERM=$TERM=/home/user/scripts`
- C. `export PATH=$PATH:/home/user/scripts`
- D. `export alias /home/user/scripts='bin'`

**Answer:** C

**Explanation:**

In Linux, the ability to execute a program by typing only its name depends on whether the directory containing the executable is included in the user's PATH environment variable. The PATH variable defines a colon-separated list of directories that the shell searches when a command is entered.

Option C, `export PATH=$PATH:/home/user/scripts`, correctly appends the /home/user/scripts directory to the existing PATH variable. Once this command is executed, any executable script located in that directory can be run simply by typing its filename, provided the script has execute permissions. This behavior is explicitly covered in the Linux+ V8 objectives related to environment variables and shell configuration.

The other options are incorrect. Option A incorrectly attempts to redefine the SHELL variable and uses invalid syntax. Option B modifies the TERM variable, which controls terminal type and has nothing to do with command execution. Option D attempts to create an alias using invalid syntax and would not affect command lookup behavior.

Linux+ V8 documentation emphasizes modifying the PATH variable as the standard and recommended method for simplifying script execution. This approach is commonly used by developers and administrators who frequently run custom scripts.

Therefore, the correct answer is C.

**NEW QUESTION 36**

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.c1
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 unconfined\_u context.
- B. The administrator forgot to restart the SSHD after creating the authorized\_keys 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.

**NEW QUESTION 38**

An administrator attempts to install updates on a Linux system but receives error messages regarding a specific repository. Which of the following commands should the administrator use to verify that the repository is installed and enabled?

- A. yum repo-pkgs
- B. yum list installed repos
- C. yum reposync available
- D. yum repolist all

**Answer: D**

**Explanation:**

Package management troubleshooting is an important skill in Linux+ V8, especially on RPM-based distributions that use yum or dnf. When update errors reference a repository, the administrator must verify whether the repository exists and whether it is enabled.

The command yum repolist all displays all configured repositories, including those that are enabled, disabled, or temporarily unavailable. This makes it the most effective command for diagnosing repository-related issues. It allows administrators to quickly confirm the repository's status and take corrective action, such as enabling it or fixing configuration errors.

The other options are incorrect. yum repo-pkgs manages packages within a repository but does not list repository status. yum list installed repos is not a valid yum command. yum reposync is used to mirror repositories locally and is not intended for verification.

Linux+ V8 documentation highlights yum repolist all as the standard command for repository inspection and troubleshooting. Therefore, the correct answer is D. yum repolist all.

**NEW QUESTION 43**

A Linux administrator is testing a web application on a laboratory service and needs to temporarily allow DNS and HTTP/HTTPS traffic from the internal network. Which of the following commands will accomplish this task?

- A. firewallld -- add-service=dns, http,https -- zone=internal
- B. iptables -- enable-service='dns|http|https' -- zone=internal
- C. firewall-cmd --add-service={dns, http, https} --zone=internal
- D. systemctl mask firewallld --for={dns, http, https} --zone=internal

**Answer: C**

**Explanation:**

Comprehensive and Detailed Explanation From Exact Extract:

The correct way to temporarily allow specific services in a particular zone with firewalld is to use firewall-cmd --add-service=service --zone=zone. Multiple services can be specified in curly braces and separated by commas. The correct syntax is:

bash CopyEdit

```
firewall-cmd --add-service={dns,http,https} --zone=internal
```

This command will allow DNS (port 53), HTTP (port 80), and HTTPS (port 443) through the firewall for the "internal" zone temporarily (for the current runtime session).

Other options:

- \* A. The command syntax is incorrect; firewalld is a service, not a command-line tool.
- \* B. iptables does not use the --enable-service flag, nor does it have zones in this way.
- \* D. systemctl mask disables services, and the rest of the command is invalid.

Reference:

CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 9: "Networking", Section: "Managing Firewalls with firewalld"

CompTIA Linux+ XK0-006 Objectives, Domain 2.0: Networking

=====

#### NEW QUESTION 48

Which of the following Ansible components contains a list of hosts and host groups?

- A. Fact
- B. Inventory
- C. Playbook
- D. Collection

**Answer: B**

#### Explanation:

Ansible architecture and core components are part of the Automation, Orchestration, and Scripting domain in CompTIA Linux+ V8. Among these components, the inventory plays a foundational role in defining the infrastructure Ansible manages.

An Ansible inventory is a file (or set of files) that contains a list of managed hosts and optionally organizes them into logical groups. These hosts can be defined by IP address, fully qualified domain name (FQDN), or hostname. Inventories may be written in INI, YAML, or dynamically generated formats. Grouping hosts allows administrators to apply configurations, roles, and tasks to multiple systems simultaneously.

Option B, Inventory, is correct because it explicitly defines which systems Ansible will target. Without an inventory, Ansible does not know where to execute tasks. Linux+ V8 documentation emphasizes inventories as the starting point for all Ansible operations.

The other options are incorrect. Facts are system variables automatically collected by Ansible about managed hosts, such as OS version or IP address. Playbooks define what actions to perform but rely on the inventory to know where to perform them. Collections are distribution units that package roles, modules, and plugins, not host definitions.

Therefore, the correct answer is B. Inventory.

#### NEW QUESTION 50

A systems administrator wants to review the logs from an Apache 2 error.log file in real time and save the information to another file for later review. Which of the following commands should the administrator use?

- A. tail -f /var/log/apache2/error.log > logfile.txt
- B. tail -f /var/log/apache2/error.log | logfile.txt
- C. tail -f /var/log/apache2/error.log >> logfile.txt
- D. tail -f /var/log/apache2/error.log | tee logfile.txt

**Answer: D**

#### Explanation:

Log monitoring is a common troubleshooting task in Linux system administration, and Linux+ V8 covers command-line tools for real-time log analysis. The requirement in this scenario is twofold: view log entries as they occur and simultaneously save them to another file.

The command tail -f /var/log/apache2/error.log | tee logfile.txt fulfills both requirements. The tail -f command follows the log file in real time, displaying new entries as they are written. The pipe (|) sends this output to the tee command, which writes the data to logfile.txt while also displaying it on standard output.

The other options are insufficient. Option A redirects output to a file but prevents real-time viewing. Option C appends output but still suppresses terminal display.

Option B is syntactically invalid and does not use a proper command for writing output.

Linux+ V8 documentation specifically references tee as a useful utility for duplicating command output streams. This makes option D the correct and most effective solution.

#### NEW QUESTION 55

Which of the following describes the method of consolidating system events to a single location?

- A. Log aggregation
- B. Health checks
- C. Webhooks
- D. Threshold monitoring

**Answer: A**

#### Explanation:

Comprehensive and Detailed 250 to 350 words of Explanation From Linux+ V8 documents:

Consolidating system events from multiple sources into a single, centralized location is a key concept in Linux system administration and is explicitly covered under logging and monitoring topics in the CompTIA Linux+ V8 objectives. This method is known as log aggregation, making option A the correct answer.

Log aggregation refers to the practice of collecting logs generated by operating systems, services, applications, and network devices and storing them in a centralized repository. In Linux environments, logs may originate from systemd-journald, syslog, application-specific log files, containers, and cloud-based workloads. Aggregating these logs allows administrators to analyze events more efficiently, correlate issues across systems, and improve troubleshooting, auditing, and security monitoring.

Linux+ V8 documentation emphasizes centralized logging as a best practice in environments with multiple servers. Without log aggregation, administrators would need to log in to each system individually to inspect logs, which is inefficient and error-prone. Centralized solutions such as syslog servers, ELK/EFK stacks, and SIEM platforms enable real-time analysis, long-term retention, and alerting based on log data.

The other options do not describe log consolidation. Health checks are used to verify whether services or systems are operational but do not collect or store event data. Webhooks are HTTP-based callbacks used for event-driven automation and notifications, not for storing logs. Threshold monitoring involves generating alerts when metrics exceed defined limits, such as CPU or memory usage, but it does not centralize system event records.

Linux+ V8 stresses that effective log aggregation improves incident response, supports compliance requirements, and enhances system visibility. It is especially important for detecting security incidents, diagnosing failures, and performing root-cause analysis across distributed systems.

**NEW QUESTION 59**

A systems administrator receives reports about connection issues to a secure web server. Given the following firewall and web server outputs:

Firewall output:

Status: active

To Action From

443/tcp DENY Anywhere

443/tcp (v6) DENY Anywhere (v6)

Web server output:

tcp LISTEN 0 4096 \*:443 :

Which of the following commands best resolves this issue?

- A. ufw disable
- B. ufw allow 80/tcp
- C. ufw delete deny https/tcp
- D. ufw allow 4096/tcp

**Answer: C**

**Explanation:**

This scenario involves firewall configuration and service accessibility, which falls under the Security domain of the CompTIA Linux+ V8 objectives. The key to resolving this issue is interpreting both the firewall output and the web server status correctly.

The web server output shows that the service is actively listening on TCP port 443, which is the standard port for HTTPS (secure web traffic). The line tcp LISTEN 0 4096 \*:443 \*: confirms that the web server is running properly and is ready to accept incoming connections on port 443 from any interface. This indicates that the problem is not with the web server configuration itself.

However, the firewall output clearly shows that incoming connections to port 443 are being blocked. The rules 443/tcp DENY Anywhere and 443/tcp (v6) DENY Anywhere (v6) indicate that the Uncomplicated Firewall (UFW) is explicitly denying HTTPS traffic for both IPv4 and IPv6. As a result, external clients cannot establish a secure connection to the server, even though the service is running correctly.

To resolve this issue securely and correctly, the administrator must remove the firewall rule that denies HTTPS traffic. Option C, ufw delete deny https/tcp, directly removes the blocking rule while preserving the rest of the firewall configuration. This aligns with Linux+ best practices, which emphasize making precise firewall changes rather than disabling security controls entirely.

The other options are incorrect. Option A, ufw disable, would completely turn off the firewall, creating a significant security risk. Option B, ufw allow 80/tcp, only opens HTTP traffic on port 80 and does not resolve HTTPS connectivity issues. Option D, ufw allow 4096/tcp, incorrectly attempts to open an internal socket backlog value rather than a valid service port.

Therefore, the correct and most secure solution is C.

**NEW QUESTION 60**

A technician wants to temporarily use a Linux virtual machine as a router for the network segment 10.10.204.0/24. Which of the following commands should the technician issue? (Select three).

- A. echo "1" > /proc/sys/net/ipv4/ip\_forward
- B. iptables -A FORWARD -j ACCEPT
- C. iptables -A PREROUTING -j ACCEPT
- D. iptables -t nat -s 10.10.204.0/24 -p tcp -A PREROUTING -j MASQUERADE
- E. echo "0" > /proc/sys/net/ipv4/ip\_forward
- F. echo "1" > /proc/net/tcp
- G. iptables -t nat -s 10.10.204.0/24 -A POSTROUTING -j MASQUERADE

**Answer: ABG**

**Explanation:**

Comprehensive and Detailed Explanation From Exact Extract:

To temporarily configure a Linux virtual machine as a router, the technician must enable IP forwarding and set up iptables rules to allow and masquerade traffic:

\* A. echo "1">/proc/sys/net/ipv4/ip\_forward: Enables IPv4 forwarding in the Linux kernel, allowing the VM to forward packets between interfaces.

\* B. iptables -A FORWARD -j ACCEPT: Adds a rule to the iptables firewall to accept all forwarded packets (allows traffic to be routed).

\* G. iptables -t nat -s 10.10.204.0/24 -A POSTROUTING -j MASQUERADE: Sets up network address translation (NAT) for outgoing packets from the 10.10.204.0/24 subnet, masquerading them as if they are coming from the VM's external IP.

Other options:

\* C.andH.are not relevant for routing/NAT in this context (PREROUTING is generally used for DNAT, not for standard source NAT).

\* D.is syntactically incorrect and mixes PREROUTING with MASQUERADE, which is not the proper combination for SNAT.

\* E.disables forwarding.

\* F.is not related to IP forwarding.

[Reference:, CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 9: "Networking", Section: "Configuring Linux as a Router", CompTIA Linux+ XK0-006 Objectives: Domain 2.0 – Networking, Official CompTIA Linux+ Cert Guide, Chapter 12: "Firewall and NAT configuration", ]

**NEW QUESTION 64**

A Linux software developer wants to use AI to optimize source code used in a commercial product. Which of the following steps should the developer take first?

- A. Research which available AI chatbots are best at optimizing source code.
- B. Verify that the company has a policy governing the use of AI in software development.
- C. Install a private LLM to use on the internal network for source code optimization.
- D. Use open-source LLMs that undergo regular security reviews by the community.

**Answer: B**

**Explanation:**

Linux+ V8 emphasizes security, compliance, and governance when introducing new automation technologies, including AI. Before using AI tools to optimize commercial source code, the developer must ensure that such usage complies with organizational policies.

Option B is correct because verifying company policy is the first and most critical step. AI tools may introduce risks such as intellectual property leakage, licensing conflicts, or regulatory violations. Many organizations restrict how source code can be shared with external systems, including AI services.

The other options are premature. Selecting tools or deploying models should only occur after policy approval. Linux+ V8 highlights governance-first approaches when adopting automation technologies.

Therefore, the correct answer is B.

#### NEW QUESTION 68

A systems administrator receives reports from users who are having issues while trying to modify newly created files in a shared directory. The administrator sees the following outputs:

```
[student3@hostname share]$ ls -ld /share
drwxrwxr-x. 5 userdata users 56 Jul 9 16:31 /share
[student3@hostname share]$ ls -l
total 4
drwxrwxr-x. 2 student users 6 Jul 9 16:28 originaldata
drwxrwxr-x. 2 student users 6 Jul 9 16:28 originalfile
-rw-rw-r--. 1 student2 student2 0 Jul 9 16:31 newfile2
drwxrwxr-x. 2 student2 student2 6 Jul 9 16:31 mynewdir
-rw-rw-r--. 1 student3 student3 0 Jul 9 16:33 newfile

[student3@hostname share]$ echo "content" >> newfile2
bash: newfile2: Permission denied

[student3@hostname share]$ touch mynewdir/file
touch: cannot touch 'mynewdir/file': Permission denied
```

Which of the following provides the best resolution to this issue?

- A. Adding a setuid bit to the user in the shared folder
- B. Manually changing the group of the newly created files
- C. Changing all directory contents to be writable and readable for everyone
- D. Adding a setgid bit to the group in the shared folder

**Answer:** D

#### Explanation:

This scenario involves shared directory collaboration, which is a common system management task covered in the CompTIA Linux+ V8 objectives. The key issue is that users can create files in the shared directory, but other users in the same group cannot modify those files. This behavior is directly related to group ownership inheritance.

By default, when a user creates a file or directory, it is owned by the user and assigned the user's primary group, not necessarily the group of the parent directory. As shown in the output, files inside /share are owned by different groups (student, student2, student3), which prevents other group members from modifying them, even though the parent directory is group-writable.

The correct solution is to set the setgid (set group ID) bit on the shared directory, making option D correct. When the setgid bit is applied to a directory, all newly created files and subdirectories inherit the group ownership of the parent directory, rather than the creator's primary group. This ensures consistent group ownership and allows all members of the shared group to collaborate effectively.

The other options are incorrect or poor practice. Option A (setuid) is intended for executables, not directories. Option B requires constant manual intervention and does not scale. Option C weakens security by granting write access to all users, violating the principle of least privilege.

Linux+ V8 documentation explicitly recommends using the setgid bit on shared directories to manage collaborative access securely and efficiently.

#### NEW QUESTION 71

Which of the following best describes journald?

- A. A system service that collects and stores logging data
- B. A feature that creates crash dumps in case of kernel failure
- C. A service responsible for keeping the filesystem journal
- D. A service responsible for writing audit records to a disk

**Answer:** A

#### NEW QUESTION 74

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

**NEW QUESTION 75**

.....

## Thank You for Trying Our Product

### We offer two products:

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

2nd - Questions and Answers in PDF Format

### XK0-006 Practice Exam Features:

- \* XK0-006 Questions and Answers Updated Frequently
- \* XK0-006 Practice Questions Verified by Expert Senior Certified Staff
- \* XK0-006 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* XK0-006 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The XK0-006 Practice Test Here](#)**