



CompTIA

Exam Questions XK0-006

CompTIA Linux+ Exam

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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. lspic
- 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. lspic: 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 needs to enable routing of IP packets between network interfaces. Which of the following kernel parameters should the administrator change?

- A. net.ipv4.ip_multicast
- B. net.ipv4.ip_route
- C. net.ipv4.ip_local_port_range
- D. net.ipv4.ip_forward

Answer: D

Explanation:

IP packet forwarding is a key networking function in Linux system management and is explicitly referenced in the Linux+ V8 objectives. Enabling this feature allows a Linux system to act as a router by forwarding packets between network interfaces.

The kernel parameter responsible for this behavior is net.ipv4.ip_forward. When this parameter is set to 1, the Linux kernel allows IPv4 packets to be forwarded between interfaces. By default, this setting is often disabled on non-routing systems for security reasons.

The parameter can be modified temporarily using the sysctl command or permanently by editing /etc/sysctl.conf or files under /etc/sysctl.d/. Linux+ V8 documentation highlights this parameter as essential for configuring routing, NAT, and firewall-based gateway systems.

The other options are incorrect. net.ipv4.ip_multicast controls multicast behavior, not packet forwarding. net.ipv4.ip_route is not a valid kernel parameter.

net.ipv4.ip_local_port_range defines the range of ephemeral ports used by outgoing connections and has no effect on routing.

Properly enabling IP forwarding is critical when configuring VPN gateways, firewalls, and network appliances. Therefore, the correct answer is D.

net.ipv4.ip_forward.

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 creating a backup copy of the `/home/` directory. Which of the following commands allows the administrator to archive and compress the directory at the same time?

- A. `cpio -o /backups/home.tar.xz /home/`
- B. `rsync -z /backups/home.tar.xz /home/`
- C. `tar -cJf /backups/home.tar.xz /home/`
- D. `dd of=/backups/home.tar.xz if=/home/`

Answer: C

Explanation:

Creating backups is a core responsibility in Linux system management, and the Linux+ V8 objectives emphasize proper use of archiving and compression tools. The `tar` utility is the standard Linux tool for creating archive files, and it also supports compression through various options.

The command `tar -cJf /backups/home.tar.xz /home/` correctly combines both archiving and compression in a single step. The `-c` option creates a new archive, `-J` specifies XZ compression, and `-f` allows the administrator to define the output file name. This results in a compressed archive of the entire `/home/` directory, which is efficient for storage and transfer.

The other options are incorrect. `cpio` is an archiving tool but does not perform compression by itself without additional commands or pipelines. `rsync -z` compresses data during transfer but does not create an archive file. The `dd` command performs low-level copying of raw data and is not suitable for directory-based backups.

Linux+ V8 documentation highlights `tar` as the preferred utility for filesystem backups due to its flexibility, reliability, and support for multiple compression algorithms. Therefore, the correct answer is C.

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

On a Kubernetes cluster, which of the following resources should be created in order to expose a port so it is publicly accessible on the internet?

- A. Deployment
- B. Network
- C. Service
- D. Pod

Answer: C

Explanation:

Container orchestration concepts are part of the Automation and Orchestration domain in Linux+ V8. In Kubernetes, workloads run inside Pods, but Pods are not directly accessible from outside the cluster.

To expose an application externally, a Service resource must be created. Services provide a stable network endpoint and can be configured as NodePort, LoadBalancer, or ClusterIP. Public exposure is typically achieved using NodePort or LoadBalancer types.

Option C, Service, is correct. Deployments manage Pods, but they do not handle networking exposure. Pods represent running containers but lack external accessibility by default. "Network" is not a valid Kubernetes resource type.

Linux+ V8 documentation highlights Services as the mechanism for exposing containerized applications. Therefore, the correct answer is C.

NEW QUESTION 8

Which of the following utilities supports the automation of security compliance and vulnerability management?

- A. SELinux
- B. Nmap
- C. AIDE
- D. OpenSCAP

Answer: D

Explanation:

Security compliance and vulnerability management are critical components of Linux system administration, and CompTIA Linux+ V8 places strong emphasis on automated security assessment tools. OpenSCAP is specifically designed to address these requirements.

OpenSCAP is an open-source framework that implements the Security Content Automation Protocol (SCAP), a set of standards used for automated vulnerability scanning, configuration compliance checking, and security auditing. It allows administrators to assess Linux systems against established security baselines such as CIS benchmarks, DISA STIGs, and organizational security policies. This makes OpenSCAP the most appropriate tool for automating both compliance and vulnerability management.

The other options serve different security-related purposes but do not fulfill the automation requirement. SELinux is a mandatory access control system that enforces security policies at runtime but does not perform compliance scanning or vulnerability assessments. Nmap is a network scanning and discovery tool used to identify open ports and services, not compliance automation. AIDE (Advanced Intrusion Detection Environment) is a file integrity monitoring tool that detects unauthorized file changes but does not evaluate overall system compliance.

Linux+ V8 documentation highlights OpenSCAP as a tool used to automate security audits, generate compliance reports, and integrate with configuration management workflows. Its ability to standardize security checks across multiple systems makes it essential in enterprise and regulated environments.

Therefore, the correct answer is D. OpenSCAP.

NEW QUESTION 9

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 needs to verify the user ID, home directory, and assigned shell for the user named "accounting." Which of the following commands should the administrator use to retrieve this information?

- A. `getent passwd accounting`
- B. `id accounting`
- C. `grep accounting /etc/shadow`
- D. `who accounting`

Answer: A

Explanation:

User account information is centrally stored in the system's account databases, and Linux+ V8 emphasizes the use of standard tools to query this data safely and consistently.

The `getent passwd accounting` command retrieves the user's entry from the `passwd` database, which may be sourced from local files or network services such as LDAP. This entry includes the username, user ID (UID), group ID (GID), home directory, and assigned login shell. Therefore, option A provides all the requested information in a single command.

Option B, `id accounting`, displays the UID and group memberships but does not show the home directory or assigned shell. Option C is incorrect because `/etc/shadow` contains password hashes and expiration data, not shell or home directory information. Option D, `who accounting`, only shows login sessions and does not provide account configuration details.

Linux+ V8 documentation highlights `getent passwd` as the preferred method for retrieving comprehensive user account information because it works across different authentication backends.

Thus, the correct answer is A.

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 13

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 17

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 19

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 tmp
[root@comptiasim ~]# ls -l /etc/lvm/archive
[root@comptiasim ~]# pvdisplay
[root@comptiasim ~]# ovs
[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 ~]# blkid
[root@comptiasim ~]# ls /
[root@comptiasim ~]# lvsdisplay
[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.

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 ~]# lvsdisplay
[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-9486-4a46-aeb6-06eaf6c47675"
TYPE="xfs"
/dev/mapper/vg01-lv01: UUID="c63883e9-ceca-45f4-9ad9-f8d8c1814e7e"
TYPE="xfs"
/dev/xvdf:          UUID="1uyvyk-Ffd0-8rvF-cYba-15Zc-EHRZ-JM3Uhm"
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
```

Select command

[illegible]

Your Partner of IT Exam

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Part 1 – Begin the recovery process Answer

vgcfsrestore 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:

vgcfsrestore 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 23

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 25

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 26

Which of the following commands should an administrator use to convert a KVM disk file to a different format?

- A. qemu-kvm
- B. qemu-ng
- C. qemu-io
- D. qemu-img

Answer: D

Explanation:

Virtualization management is part of Linux system administration and is included in Linux+ V8 objectives. KVM virtual machines commonly use disk image formats such as qcow2, raw, or vmdk. Converting between these formats is a routine administrative task.

The correct tool for disk image conversion is `qemu-img`. This utility allows administrators to create, convert, resize, and inspect virtual disk images. For example, converting a qcow2 image to raw format can be accomplished using `qemu-img convert`. This capability is explicitly referenced in Linux+ V8 documentation related to virtualization tooling.

The other options are incorrect. `qemu-kvm` refers to the hypervisor component, not disk manipulation. `qemu-ng` is not a valid QEMU utility. `qemu-io` is used for low-level I/O testing and debugging, not image format conversion.

Therefore, the correct answer is D. `qemu-img`.

NEW QUESTION 29

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 32

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 33

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. `grep a^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.

NEW QUESTION 36

A systems administrator is having issues with a third-party API endpoint. The administrator receives the following output:

```
# curl https://comptia.com/endpoint
curl: (6) Could not resolve host: comptia.com

# dig comptia.com
; <<>> <<>> comptia.com
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NXDOMAIN, id: 14031
;; flags: qr rd ra; QUERY: 1, ANSWER: 0, AUTHORITY: 1, ADDITIONAL: 1
;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 512
;; QUESTION SECTION:
;comptia.com. IN A
;; AUTHORITY SECTION:
com. 900 IN SOA a.gtld-servers.net. nstld.verisign-grs.com. 1720473015 1800 900 604800 86400
;; Query time: 159 msec
;; SERVER: 10.255.255.254#53(10.255.255.254) (UDP)
;; WHEN: Mon Jul 08 15:10:45 CST 2024
;; MSG SIZE rcvd: 117
```

Which of the following actions should the administrator take to resolve the issue?

- A. Open a secure port in the server's firewall.
- B. Request a new API endpoint from a third party.
- C. Review and fix the DNS client configuration file.
- D. Enable internet connectivity on the host.

Answer: C

NEW QUESTION 38

.....

Relate Links

100% Pass Your XK0-006 Exam with Exam Bible Prep Materials

<https://www.exambible.com/XK0-006-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>