CompTIA Linux+ XK0-006 Practice Questions

Shared by Bond on 13-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Which of the following best describes journald?

Options:

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

Explanation:

journald, part of systemd, is the core logging service in modern Linux systems and is covered under Linux+ V8 logging and monitoring objectives.

The correct description is A. systemd-journald collects, stores, and indexes logging data from the kernel, system services, and applications. Logs are stored in a structured, binary format and can be queried using journalctl. Journald supports metadata tagging, log filtering, and centralized logging integration.

Option B refers to kernel crash dump mechanisms like kdump. Option C describes filesystem journaling (such as ext4 journaling). Option D refers to auditd, which manages security audit logs.

Linux+ V8 documentation clearly distinguishes journald from other logging and auditing services. Therefore, the correct answer is A.

Question 2

Question Type: MultipleChoice

An administrator added a new disk to expand the current storage. Which of the following commands should the administrator run first to add the new disk to the LVM?

Options:

A- vgextend
B- lvextend
C- pvcreate
D- pvresize

Answer:

C

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

To add a new physical disk to LVM, the disk must first be initialized as a physical volume using the pvcreate command. This prepares the new disk for use by the LVM subsystem. After initializing with pvcreate, you would use vgextend to add the new physical volume to an existing volume group.

Other options:

A . vgextend adds a physical volume to a volume group, but you must use pvcreate first.

B . lvextend is used to increase the size of a logical volume, not to add a new disk.

D . pvresize is used to resize an existing physical volume, not to create one.

CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 7: 'Managing Storage', Section: 'Managing Logical Volumes'

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

Question 3

Question Type: MultipleChoice

Which of the following describes how a user's public key is used during SSH authentication?

Options:

A- The user's public key is used to hash the password during SSH authentication.
B- The user's public key is verified against a list of authorized keys. If it is found, the user is

allowed to log in.
C- The user's public key is used instead of a password to allow server access.
D- The user's public key is used to encrypt the communication between the client and the server.

Answer:

B

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

During SSH public key authentication, the server checks if the user's public key is present in the ~/.ssh/authorized_keys file. If the key is found, the server uses it to verify the user's identity by sending a challenge that can only be answered by the corresponding private key. This process does not involve password hashing or using the public key directly for encryption of the communication stream. Instead, the public key is simply used as a reference for authentication.

CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 11: 'Securing Linux', Section: 'SSH Key-Based Authentication'

CompTIA Linux+ XK0-006 Objectives, Domain 3.0: Security

==========

Question 4

Question Type: MultipleChoice

Which of the following commands should a Linux administrator use to determine the version of a kernel module?

Options:

A- modprobe bluetooth
B- lsmod bluetooth
C- depmod bluetooth
D- modinfo bluetooth

Answer:

D

Explanation:

Kernel module management is an important part of Linux system administration and is covered in the Linux+ V8 objectives. When an administrator needs to determine metadata about a kernel module---such as its version, author, description, license, filename, and dependencies---the correct tool is modinfo.

The command modinfo bluetooth displays detailed information about the specified kernel module, including the module version if it is defined. This makes it the correct and intended command for retrieving version details of kernel modules, whether or not the module is currently loaded.

The other options are incorrect. modprobe bluetooth is used to load or unload kernel modules and does not display version information. lsmod lists loaded modules but does not show version details and does not accept module names as arguments in that manner. depmod is used to generate module dependency information and does not provide module metadata to the administrator.

Linux+ V8 documentation specifically references modinfo as the utility for inspecting kernel module properties. This command is essential for troubleshooting driver issues, verifying compatibility, and auditing kernel components.

Therefore, the correct answer is D. modinfo bluetooth.

Question 5

Question Type: MultipleChoice

A Linux user needs to download the latest Debian image from a Docker repository. Which option best commands makes this task possible?

Options:

A- docker image init debian
B- docker image pull debian
C- docker image import debian
D- docker image save debian

Answer:

B

Explanation:

Container management and image handling are part of modern Linux automation practices covered in CompTIA Linux+ V8. Docker images are stored in container registries such as Docker Hub, and administrators commonly need to download images to deploy containers.

The correct command for downloading an image from a Docker repository is docker image pull. This command retrieves the specified image from a configured container registry and stores it locally. When no tag is specified, Docker automatically pulls the latest available version of the image. Therefore, docker image pull debian downloads the most recent Debian image from Docker Hub.

The other options are incorrect. docker image init is not a valid Docker command and does not exist in Docker's CLI. docker image import is used to create a Docker image from a tarball file, not to download an image from a repository. docker image save exports an existing local image into a tar archive and does not retrieve images from a remote registry.

Linux+ V8 documentation emphasizes understanding container image lifecycles, including pulling, tagging, and running images. Pulling images is a foundational step before container execution and automation workflows.

Therefore, the correct answer is B. docker image pull debian.

Question 6

Question Type: MultipleChoice

SIMULATION

Joe, a user, has taken a position previously held by Ann. As a systems administrator, you

need to archive all the files from Ann's home directory and extract them into Joe's home

directory.

INSTRUCTIONS

Within each tab, click on an object to form the appropriate commands. Command objects may only be used once, but the spacebar _ object may be used multiple times. Not all objects will be used.

If at any time you would like to bring back the initial state of the simulation, please click the Reset All button.

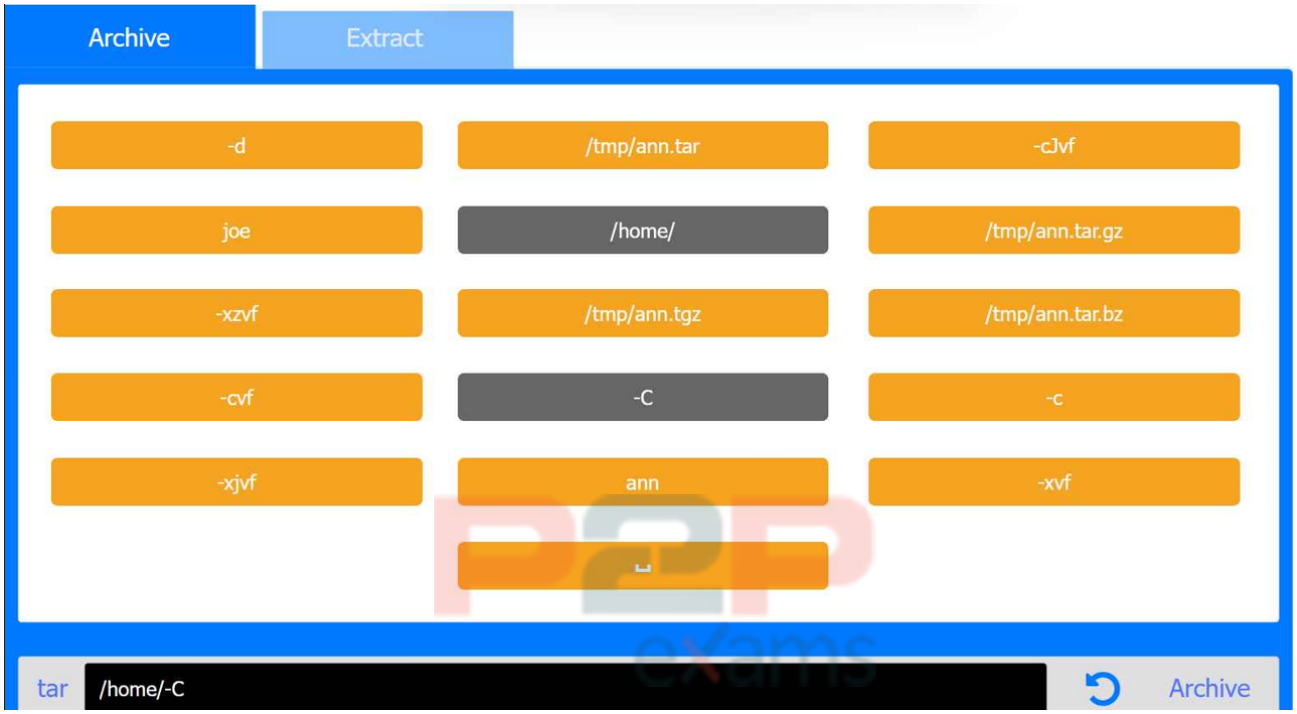

Options:

A- See the solution in Explanation below

Answer:

A

Explanation:

Archive Tab -- Create the archive from Ann's home directory

Correct Command:

tar -cvf /tmp/ann.tar -C /home/ ann

Extract Tab -- Extract the archive into Joe's home directory

Correct Command:

tar -xvf /tmp/ann.tar -C /home/ joe

This performance-based question tests file archiving and restoration using tar, a core System Management skill in the CompTIA Linux+ V8 objectives. The task requires preserving Ann's files and placing them correctly into Joe's home directory.

Archive Phase Explanation

The goal of the first step is to archive Ann's entire home directory without embedding the full path (/home/ann) inside the archive. This is accomplished using the -C option.

Command breakdown:

tar archive utility

-c create an archive

-v verbose output (optional but allowed)

-f /tmp/ann.tar specifies the archive file

-C /home/ changes directory before archiving

ann archives the ann directory only

This results in a clean archive containing Ann's files without absolute paths, which is best practice and explicitly covered in Linux+ V8 documentation.

Extract Phase Explanation

The second step extracts the archived files into Joe's home directory.

Command breakdown:

-x extract

-v verbose

-f /tmp/ann.tar specifies the archive

-C /home/joe extracts files directly into Joe's home directory

This ensures Joe receives Ann's files correctly under /home/joe/ann or directly under /home/joe depending on post-extraction handling, which matches Linux+ expectations for administrative user transitions.

Question 7

Question Type: MultipleChoice

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?

Options:

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.

Question 8

Question Type: MultipleChoice

A systems administrator manages multiple Linux servers and needs to set up a reliable and secure way to handle the complexity of managing event records on the OS and application levels. Which of the following should the administrator do?

Options:

A- Create an automated process to retrieve logs from the server by demand.
B- Implement a centralized log aggregation solution.
C- Configure daily automatic backups of logs to remote storage.
D- Deploy log rotation procedures to manage the records.

Answer:

B

Explanation:

Log management is a critical system management function highlighted in CompTIA Linux+ V8, particularly in multi-server environments. As the number of systems and applications grows, managing logs locally on each server becomes inefficient and error-prone.

The best solution is to implement a centralized log aggregation solution, making option B correct. Centralized logging collects logs from multiple systems and applications into a single, secure location. This simplifies monitoring, searching, correlation, auditing, and incident response. Common solutions include syslog servers, ELK/EFK stacks, and SIEM platforms.

Linux+ V8 documentation emphasizes centralized logging as a best practice for availability, troubleshooting, and security analysis. It enables administrators to detect patterns, investigate incidents, and maintain compliance more effectively than isolated log files.

The other options are insufficient on their own. On-demand retrieval does not scale well. Log backups protect data but do not simplify analysis. Log rotation manages disk usage but does not address distributed log complexity.

Therefore, the correct answer is B. Implement a centralized log aggregation solution.

Question 9

Question Type: MultipleChoice

A systems administrator is configuring new Linux systems and needs to enable passwordless authentication between two of the servers. Which of the following commands should the administrator use?

Options:

A- ssh-keygen -t rsa && ssh-copy-id -i ~/.ssh/id_rsa.pub john@server2
B- ssh-keyscan -t rsa && ssh-copy-id john@server2 -i ~/.ssh/key
C- ssh-agent -i rsa && ssh-copy-id ~/.ssh/key john@server2
D- ssh-add -t rsa && scp -rp ~/.ssh john@server2

Answer:

A

Explanation:

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

Passwordless authentication using SSH key pairs is a foundational security practice covered in the Security domain of CompTIA Linux+ V8. It allows administrators to securely authenticate between systems without transmitting passwords over the network, significantly reducing the risk of credential compromise.

The correct approach involves two essential steps: generating an SSH key pair and installing the public key on the remote system. Option A correctly performs both steps using best-practice commands.

The command ssh-keygen -t rsa generates an RSA public/private key pair in the user's ~/.ssh/ directory. The private key (id_rsa) remains securely on the local system, while the public key (id_rsa.pub) is intended to be shared. The second part of the command, ssh-copy-id -i ~/.ssh/id_rsa.pub john@server2, securely copies the public key to the remote server's ~/.ssh/authorized_keys file. This enables key-based authentication for the specified user.

The other options are incorrect or incomplete. Option B uses ssh-keyscan, which is intended for collecting host keys to populate known_hosts, not for user authentication. Option C misuses ssh-agent, which manages keys already generated and does not create or install them. Option D is insecure and incorrect because copying the entire .ssh directory risks exposing private keys and violates security best practices.

Linux+ V8 documentation emphasizes the use of ssh-keygen and ssh-copy-id as the standard, secure method for configuring passwordless SSH access. This approach ensures proper permissions, correct key placement, and minimal risk.

Question 10

Question Type: MultipleChoice

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?

Options:

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.

Question 11

Question Type: MultipleChoice

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-killer

mysqld cpuset=/ mems_allowed=0

Which option best explains why the server was offline?

Options:

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:

This scenario clearly indicates a memory exhaustion condition, which falls under the Troubleshooting domain of the CompTIA Linux+ V8 objectives. The most critical clue is the log entry stating that mysqld invoked oom-killer.

The OOM (Out-Of-Memory) killer is a Linux kernel mechanism that activates when the system runs critically low on available memory and cannot satisfy memory allocation requests. When this happens, the kernel selects a process---typically one consuming a large amount of memory---and forcibly terminates it to protect overall system stability. In this case, the MySQL daemon (mysqld) was identified as the process responsible for triggering the OOM condition.

The journalctl output explicitly confirms this behavior. Linux+ V8 documentation emphasizes that when the OOM killer is invoked, it is almost always due to physical memory exhaustion or insufficient swap space, not user intervention or application bugs alone. The additional log line showing mems_allowed=0 further supports the conclusion that the process could not allocate memory from available memory nodes.

The fact that uptime -p reports only 2 minutes of uptime strongly suggests that the system was either rebooted automatically or manually following the memory exhaustion event. Systems may reboot as part of recovery procedures after severe resource exhaustion, especially in production environments.

The other options can be ruled out. There is no indication of a user-initiated kill signal, filesystem corruption, or network connectivity issues. Network outages would not generate OOM killer messages, and filesystem errors would appear as I/O or disk-related errors in the logs.

Linux+ V8 best practices recommend addressing OOM issues by increasing system memory, tuning MySQL memory parameters, configuring swap space, or adjusting OOM scoring.

Therefore, the correct explanation is A. The process exhausted server memory.

Question 12

Question Type: MultipleChoice

Users report that a Linux system is unresponsive and simple commands take too long to complete. The Linux administrator logs in to the system and sees the following: Output 1:

10:06:29 up 235 day, 19:23, 2 users, load average: 8.71, 8.24, 7.71

Output 2:

```
Linux 6.8.0-31-generic (host) 05/10/2024 _x86_64_ (4 CPU)
```

10:07:42	AM	CPU	%usr	%nice	%sys	%iowait	%irq	%soft	%steal	%guest	%gnice	%idle
10:07:42	AM	all	65.88	0	20.54	5.65	0	7.93	0	0	0	0

Which of the following is the system experiencing?

Options:

A- High latency
B- High uptime
C- High CPU load
D- High I/O wait times

Answer:

C

Explanation:

This scenario is a classic performance troubleshooting case covered under the Troubleshooting domain of the CompTIA Linux+ V8 objectives. The key indicators to analyze are the load average values and the CPU utilization statistics.

The uptime command shows load averages of 8.71, 8.24, and 7.71 over the 1-, 5-, and 15-minute intervals. Load average represents the average number of processes that are either running on the CPU or waiting to run. On a system with 4 CPU cores, a healthy load average would typically be close to or below 4. Load averages consistently near or above 8 indicate that there are significantly more runnable processes than available CPU resources, causing processes to wait and resulting in poor system responsiveness.

The CPU output further confirms this condition. The %idle value is 0, meaning the CPU has no

idle time available. The majority of CPU time is spent in user space (65.88%) and system/kernel space (20.54%), indicating heavy computational and kernel activity. While %iowait is present at 5.65%, it is not high enough to suggest that disk I/O is the primary bottleneck.

Option C, high CPU load, best explains the symptoms. High CPU load causes commands to execute slowly because processes are competing for limited CPU time. This directly matches the observed behavior of the system being unresponsive.

The other options are incorrect. High uptime simply indicates how long the system has been running and does not cause performance issues by itself. High latency is a general term and not a specific diagnosis shown by the metrics provided. High I/O wait times would require a significantly higher %iowait value.

According to Linux+ V8 documentation, correlating load averages with CPU core count and utilization is essential for accurate performance diagnosis. Therefore, the correct answer is C. High CPU load.

Question 13

Question Type: MultipleChoice

A systems administrator is writing a script to analyze the number of files in the directory /opt/application/home/. Which option best commands should the administrator use in conjunction with ls -l | to count the files?

Options:

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

Answer:

D

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.

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