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
Which of the following is a protocol for accessing distributed directory services containing a hierarchy of users, groups, machines, and organizational units?

A. SMB
B. TLS
C. LDAP
D. KRB-5

Answer: C

Explanation:
Directory services are a key part of enterprise Linux environments and are covered under the Security domain in Linux+ V8. The Lightweight Directory Access Protocol (LDAP) is specifically designed to access and manage distributed directory information.
LDAP directories store structured, hierarchical data such as users, groups, computers, and organizational units. Linux systems commonly use LDAP for centralized authentication, authorization, and identity management. LDAP is also the foundation for services like Active Directory and FreeIPA.
The other options are incorrect. SMB is a file and printer sharing protocol. TLS is an encryption protocol used to secure communications. Kerberos (KRB-5) is an authentication protocol often used alongside LDAP but does not store directory information itself.
Linux+ V8 documentation highlights LDAP as the primary protocol for directory-based identity services. Therefore, the correct answer is C.

NEW QUESTION 3
An administrator updates the network configuration on a server but wants to ensure the change will not cause an outage if something goes wrong. Which of the following commands allows the administrator to accomplish this goal?

A. netplan try
B. netplan rebind
C. netplan ip
D. netplan apply

Answer: A

Explanation:
Network configuration changes can cause immediate loss of connectivity if applied incorrectly. Linux+ V8 emphasizes safe configuration practices, particularly when managing remote systems.
The netplan try command applies network configuration changes temporarily and prompts the administrator to confirm them within a timeout period. If the administrator does not confirm, Netplan automatically rolls back to the previous working configuration. This prevents accidental outages caused by misconfigured network settings.
The netplan apply command makes changes permanent immediately and does not provide rollback protection. The other options are not valid Netplan commands.
Linux+ V8 documentation explicitly references netplan try as a safe testing mechanism. Therefore, the correct answer is A.

NEW QUESTION 4
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 5
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 6
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 7
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 anHTTP-based callback mechanismthat allows one system to notify another when a specific event occurs.
OptionDcorrectly 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. OptionAconfuses webhooks with authentication mechanisms. OptionBincorrectly associates webhooks with SNMP, which is a separate protocol. OptionCis 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 isD.

NEW QUESTION 8
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 9
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, aServiceresource 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.
OptionC, 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 isC.

NEW QUESTION 10
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,playbooksare most commonly associated with configuration management tools such asAnsible, which is explicitly referenced in the CompTIA Linux+ V8 objectives. Playbooks are written in YAML and are designed to definea series of tasks, configurations, and desired system statesthat should be applied to one or more Linux systems in a repeatable and automated manner.
A primary use case for playbooks isapplication 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 optionA, 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. OptionBrefers 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. OptionCdescribes container security information, which is typically managed through container runtime configurations, secrets, or security policies rather than playbooks. OptionDrefers 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 isA, as it best describes the practical and documented use case for playbooks in a Linux system.

NEW QUESTION 10
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?

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 toimplement a centralized log aggregation solution, making optionBcorrect. 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 isB. Implement a centralized log aggregation solution.

NEW QUESTION 13
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 as104.18.16.29. 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 18
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 thepasswd 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, optionAprovides all the requested information in a single command.
OptionB, id accounting, displays the UID and group memberships but does not show the home directory or assigned shell. OptionCis incorrect because /etc/shadow contains password hashes and expiration data, not shell or home directory information. OptionD, 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 isA.

NEW QUESTION 21
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 theSecuritydomain of CompTIA Linux+ V8. When a user goes on extended leave, the goal is toprevent unauthorized accesswhile preserving the user??s data and account for future use.
The most effective approach is todisable authentication and interactive login accesswithout deleting the account. OptionD, 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.
OptionF, 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. OptionA(immutable files) does not prevent account access and may interfere with system operations. OptionBdefeats the purpose of securing an inactive account. OptionCdeletes user data, which is unnecessary and risky. OptionEhas no security effect, as filesystem timestamps do not control access.
Linux+ V8 stresses that secure account management should bereversible, 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 23
A Linux administrator is making changes to local files that are part of a Git repository. The administrator needs to retrieve changes from the remote Git repository. Which of the following commands should the administrator use to save the local modifications for later review?

A. git stash
B. git pull
C. git merge
D. git fetch

Answer: A

Explanation:
In Git-based workflows, especially those used in DevOps environments, it is common for administrators to have uncommitted local changes while needing to retrieve updates from a remote repository. Linux+ V8 emphasizes understanding how to safely manage local modifications during synchronization operations.
The command git stash is specifically designed for this scenario. It temporarily saves (or ??stashes??) local changes in a stack-like structure and reverts the working directory to a clean state that matches the current HEAD. This allows the administrator to perform operations such as git pull without conflicts. Later, the stashed changes can be reapplied using git stash apply or git stash pop.
The other options are incorrect. git pull retrieves and merges remote changes but will fail or cause conflicts if local modifications exist. git merge combines branches and does not save uncommitted changes. git fetch downloads remote references but does not address local working directory changes.
Linux+ V8 documentation highlights git stash as a safe and reversible way to protect local work during repository updates. Therefore, the correct answer is A.

NEW QUESTION 26
A Linux administrator receives reports that an application hosted in a system is not completing tasks in the allocated time. The administrator connects to the system and obtains the following details:

```
# uptime
12:47:43 up 22:17, 2 users, load average: 7.75, 5.72, 5.17

# nproc
4

# vmstat -w 1 3
[...]
r b swpd free    buffcachesisobibo in   cs   us syidwa stgu
8 0671563760348103671476 0  0  0  040901386100 0  0  0  0  0
8 0671563760348103671476 0  0  0  040761389100 0  0  0  0  0
8 0671563760348103671476 0  0  0  040761389100 0  0  0  0  0

# free -h
       total  used  free shared buff/cache available
Mem:   3.8Gi 334Mi 3.6Gi   20Mi       70Mi     3.5Gi
Swap:  7.8Gi  65Mi 7.8Gi
```

Which of the following actions can the administrator take to help speed up the jobs?

A. Increase the amount of free memory available to the system.
B. Increase the amount of CPU resources available to the system.
C. Increase the amount of swap space available to the system.
D. Increase the amount of disks available to the system.

Answer: B

Explanation:
This scenario represents a classic CPU-bound performance issue, which is covered under the Troubleshooting domain of CompTIA Linux+ V8. The most important indicator is the load average compared to the number of available CPU cores.
The system has 4 CPU cores, as shown by nproc, but the load averages are consistently above 5, with a peak of 7.75. Load average reflects the number of processes either actively running on the CPU or waiting for CPU time. When the load average exceeds the number of CPU cores for extended periods, it indicates CPU contention. Processes must wait longer to be scheduled, resulting in delayed task completion.
The memory statistics confirm that memory is not the bottleneck. free -h shows over 3.5 GiB of available memory, and swap usage is minimal. Additionally, vmstat shows no significant swap-in or swap-out activity and low I/O wait, ruling out memory pressure and disk bottlenecks.
Increasing swap space would not help because the system is not memory constrained. Adding more disks would not address CPU scheduling delays. Increasing free memory is unnecessary because sufficient memory is already available.
Linux+ V8 documentation emphasizes correlating load average with CPU core count to diagnose CPU saturation. The most effective way to speed up job execution in this case is to increase CPU resources, such as adding more vCPUs, moving the workload to a more powerful system, or distributing the workload across multiple systems.
Therefore, the correct answer is B. Increase the amount of CPU resources available to the system.

NEW QUESTION 28
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 31
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 36
A DevOps engineer needs to create a local Git repository. Which of the following commands should the engineer use?

A. git init
B. git clone
C. git config
D. git add

Answer: A

Explanation:
Version control is a core DevOps practice, and CompTIA Linux+ V8 includes Git fundamentals as part of automation and orchestration objectives. To create a new local Git repository, the correct command is git init.
The git init command initializes a new Git repository in the current directory by creating a hidden .git directory. This directory contains all the metadata required for version control, including commit history, branches, configuration settings, and object storage. After running git init, the directory becomes a fully functional local repository ready to track files and commits.
The other options do not create a new repository. git clone is used to copy an existing remote repository to a local system, not to create a new one. git config is used to set Git configuration values such as username, email, or default editor. git add stages files for commit but only works after a repository has already been

initialized.
Linux+ V8 documentation highlights git init as the foundational command for starting version control in new projects. This command is frequently used in DevOps workflows when creating infrastructure-as-code repositories, automation scripts, or application source trees from scratch.
By initializing a local repository, engineers can begin tracking changes, collaborating with others, and integrating Git into CI/CD pipelines. Therefore, the correct answer is A. git init.

NEW QUESTION 41
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 46
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.

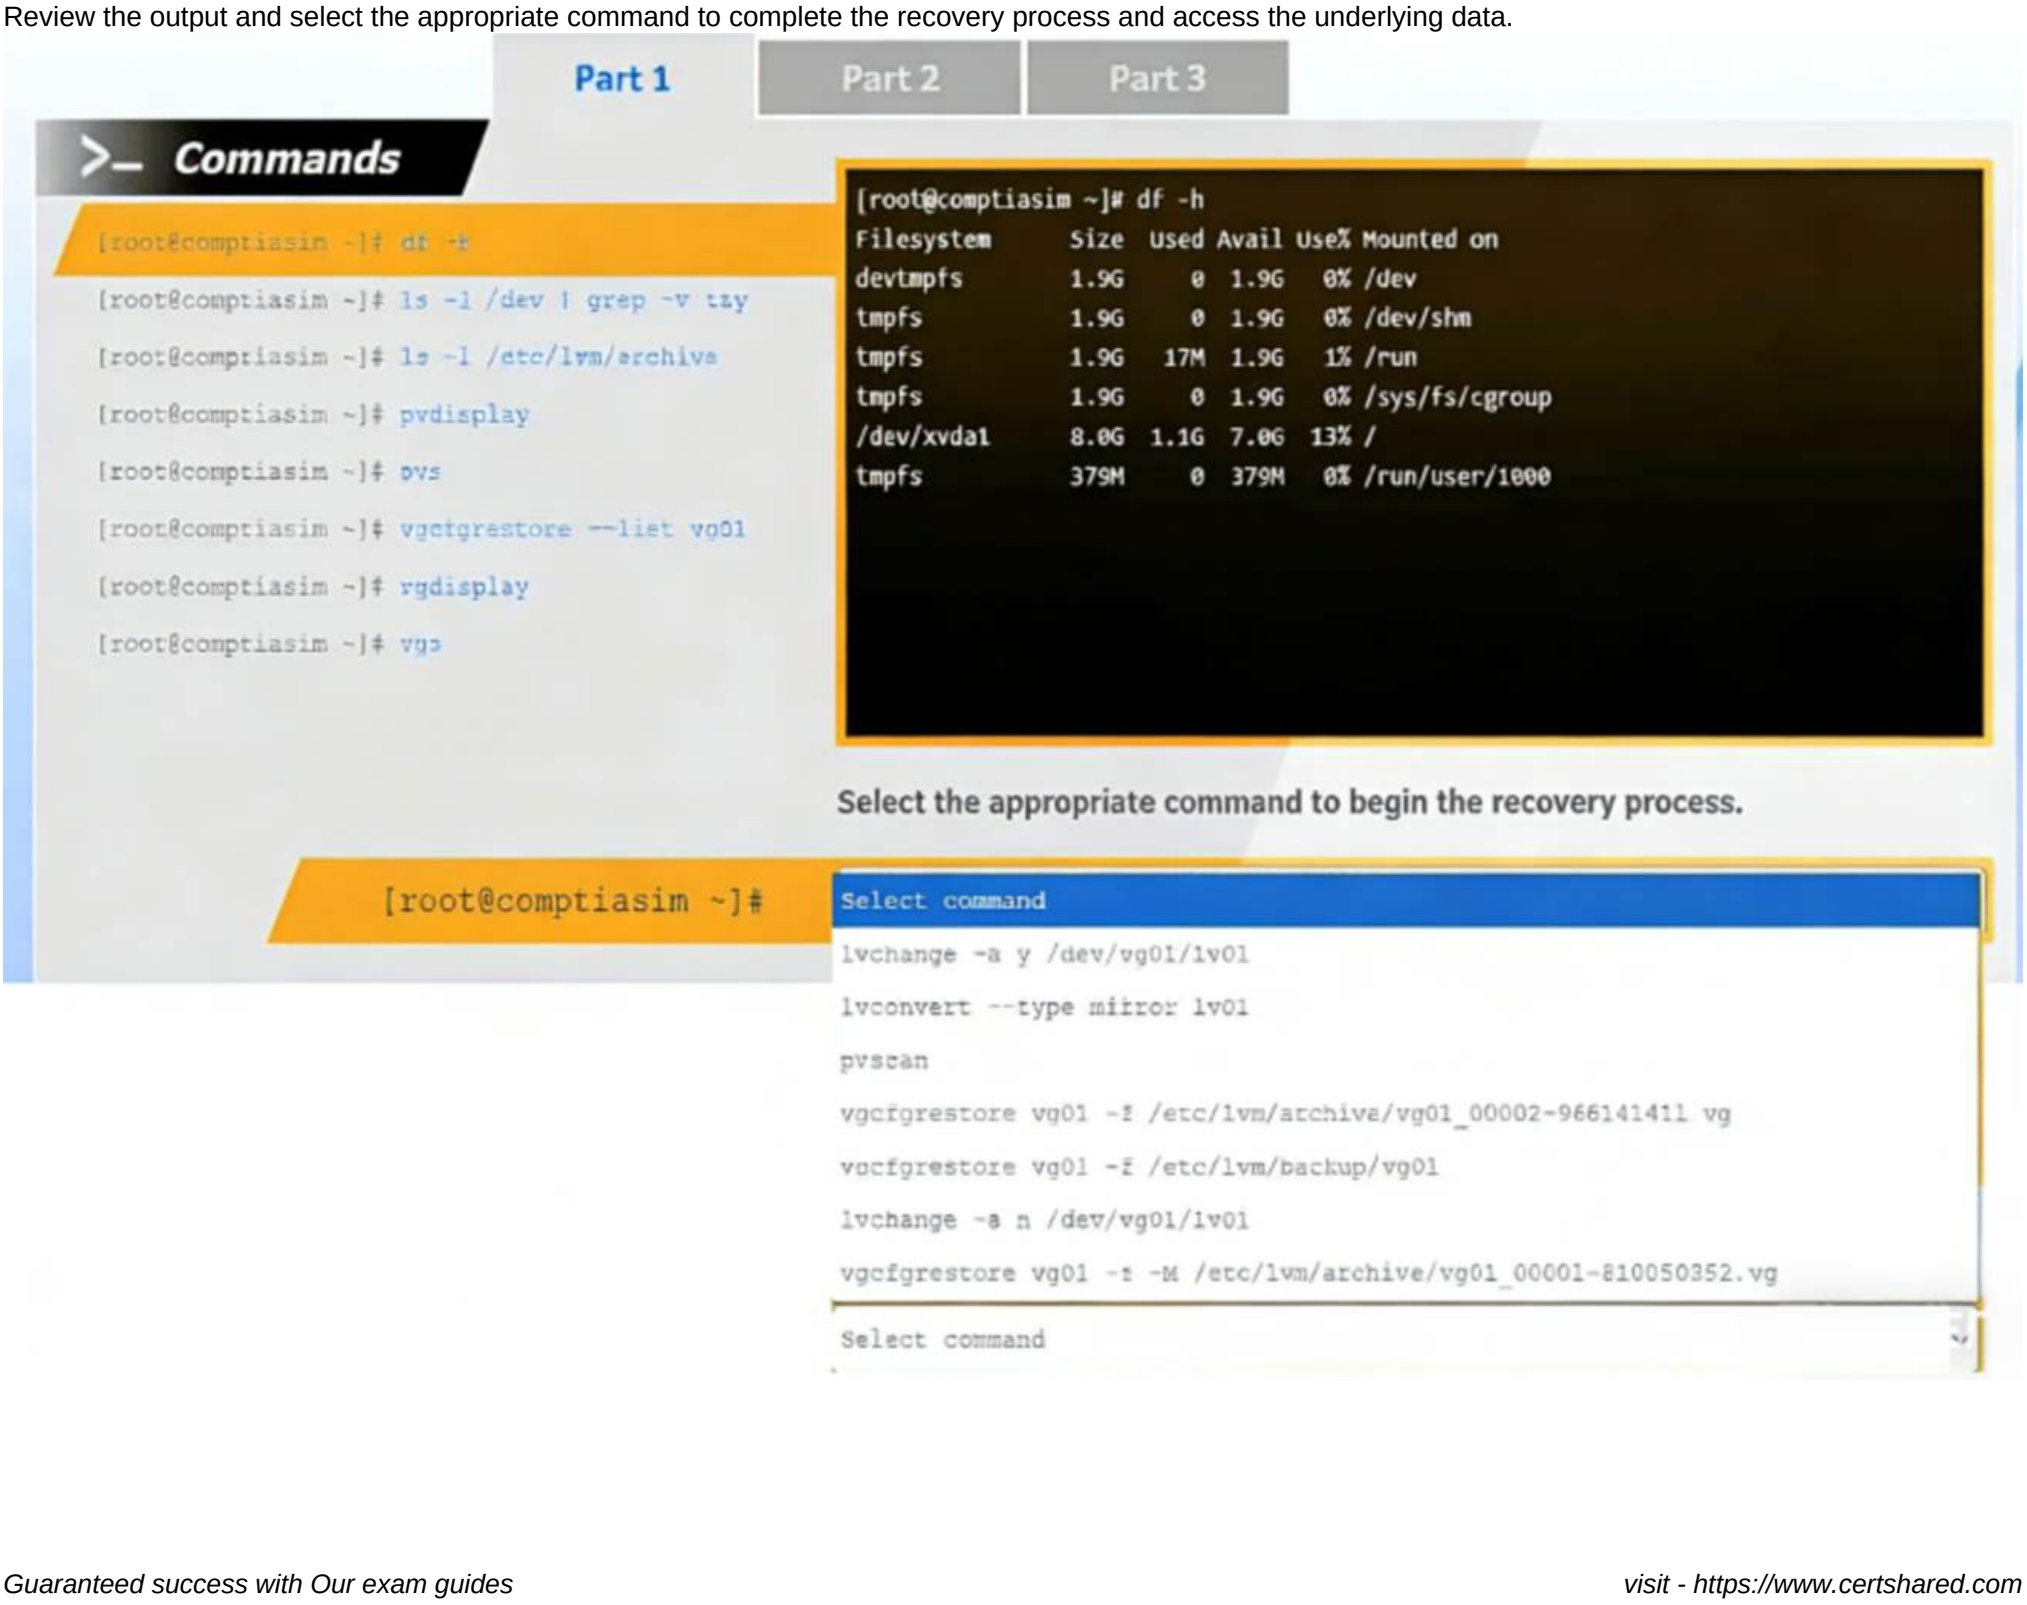

Part 2 | **Part 2** | Part 2

>_ Commands

```
[root@comptiasim ~]# blkid
[root@comptiasim ~]# dmesg | tail -20
[root@comptiasim ~]# blbik
[root@comptiasim ~]# ls /
[root@comptiasim ~]# lvdisplay
[root@comptiasim ~]# lvs
[root@comptiasim ~]# lvecan
[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-lSzC-EHRZ-JN3UHm" 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/fstap
[root@comptiasim ~]# ls -l /dev/napper/
[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-9486-4a46-aeb6-06eaf6c47675"
TYPE="xfs"
/dev/mapper/vg01-lv01:   UUID="c63883e9-ceca-45f4-9ad9-f8d8c1814e7e"
TYPE="xfs"
/dev/xvdf:               UUID="1uyvyk-Ffd0-RrYF-CYba-l5zC-EHRZ UN3UHm"
TYPE="LVM2_member"
```

```
Select command
xfs_repait /dev/vg01/lv01
lvscan -a
mount -a
mount /important_data /dev/vg01/lv01
xfs_mdrestore /dev/vg01 /important_data
```

[root@comptiasim ~]# Select command

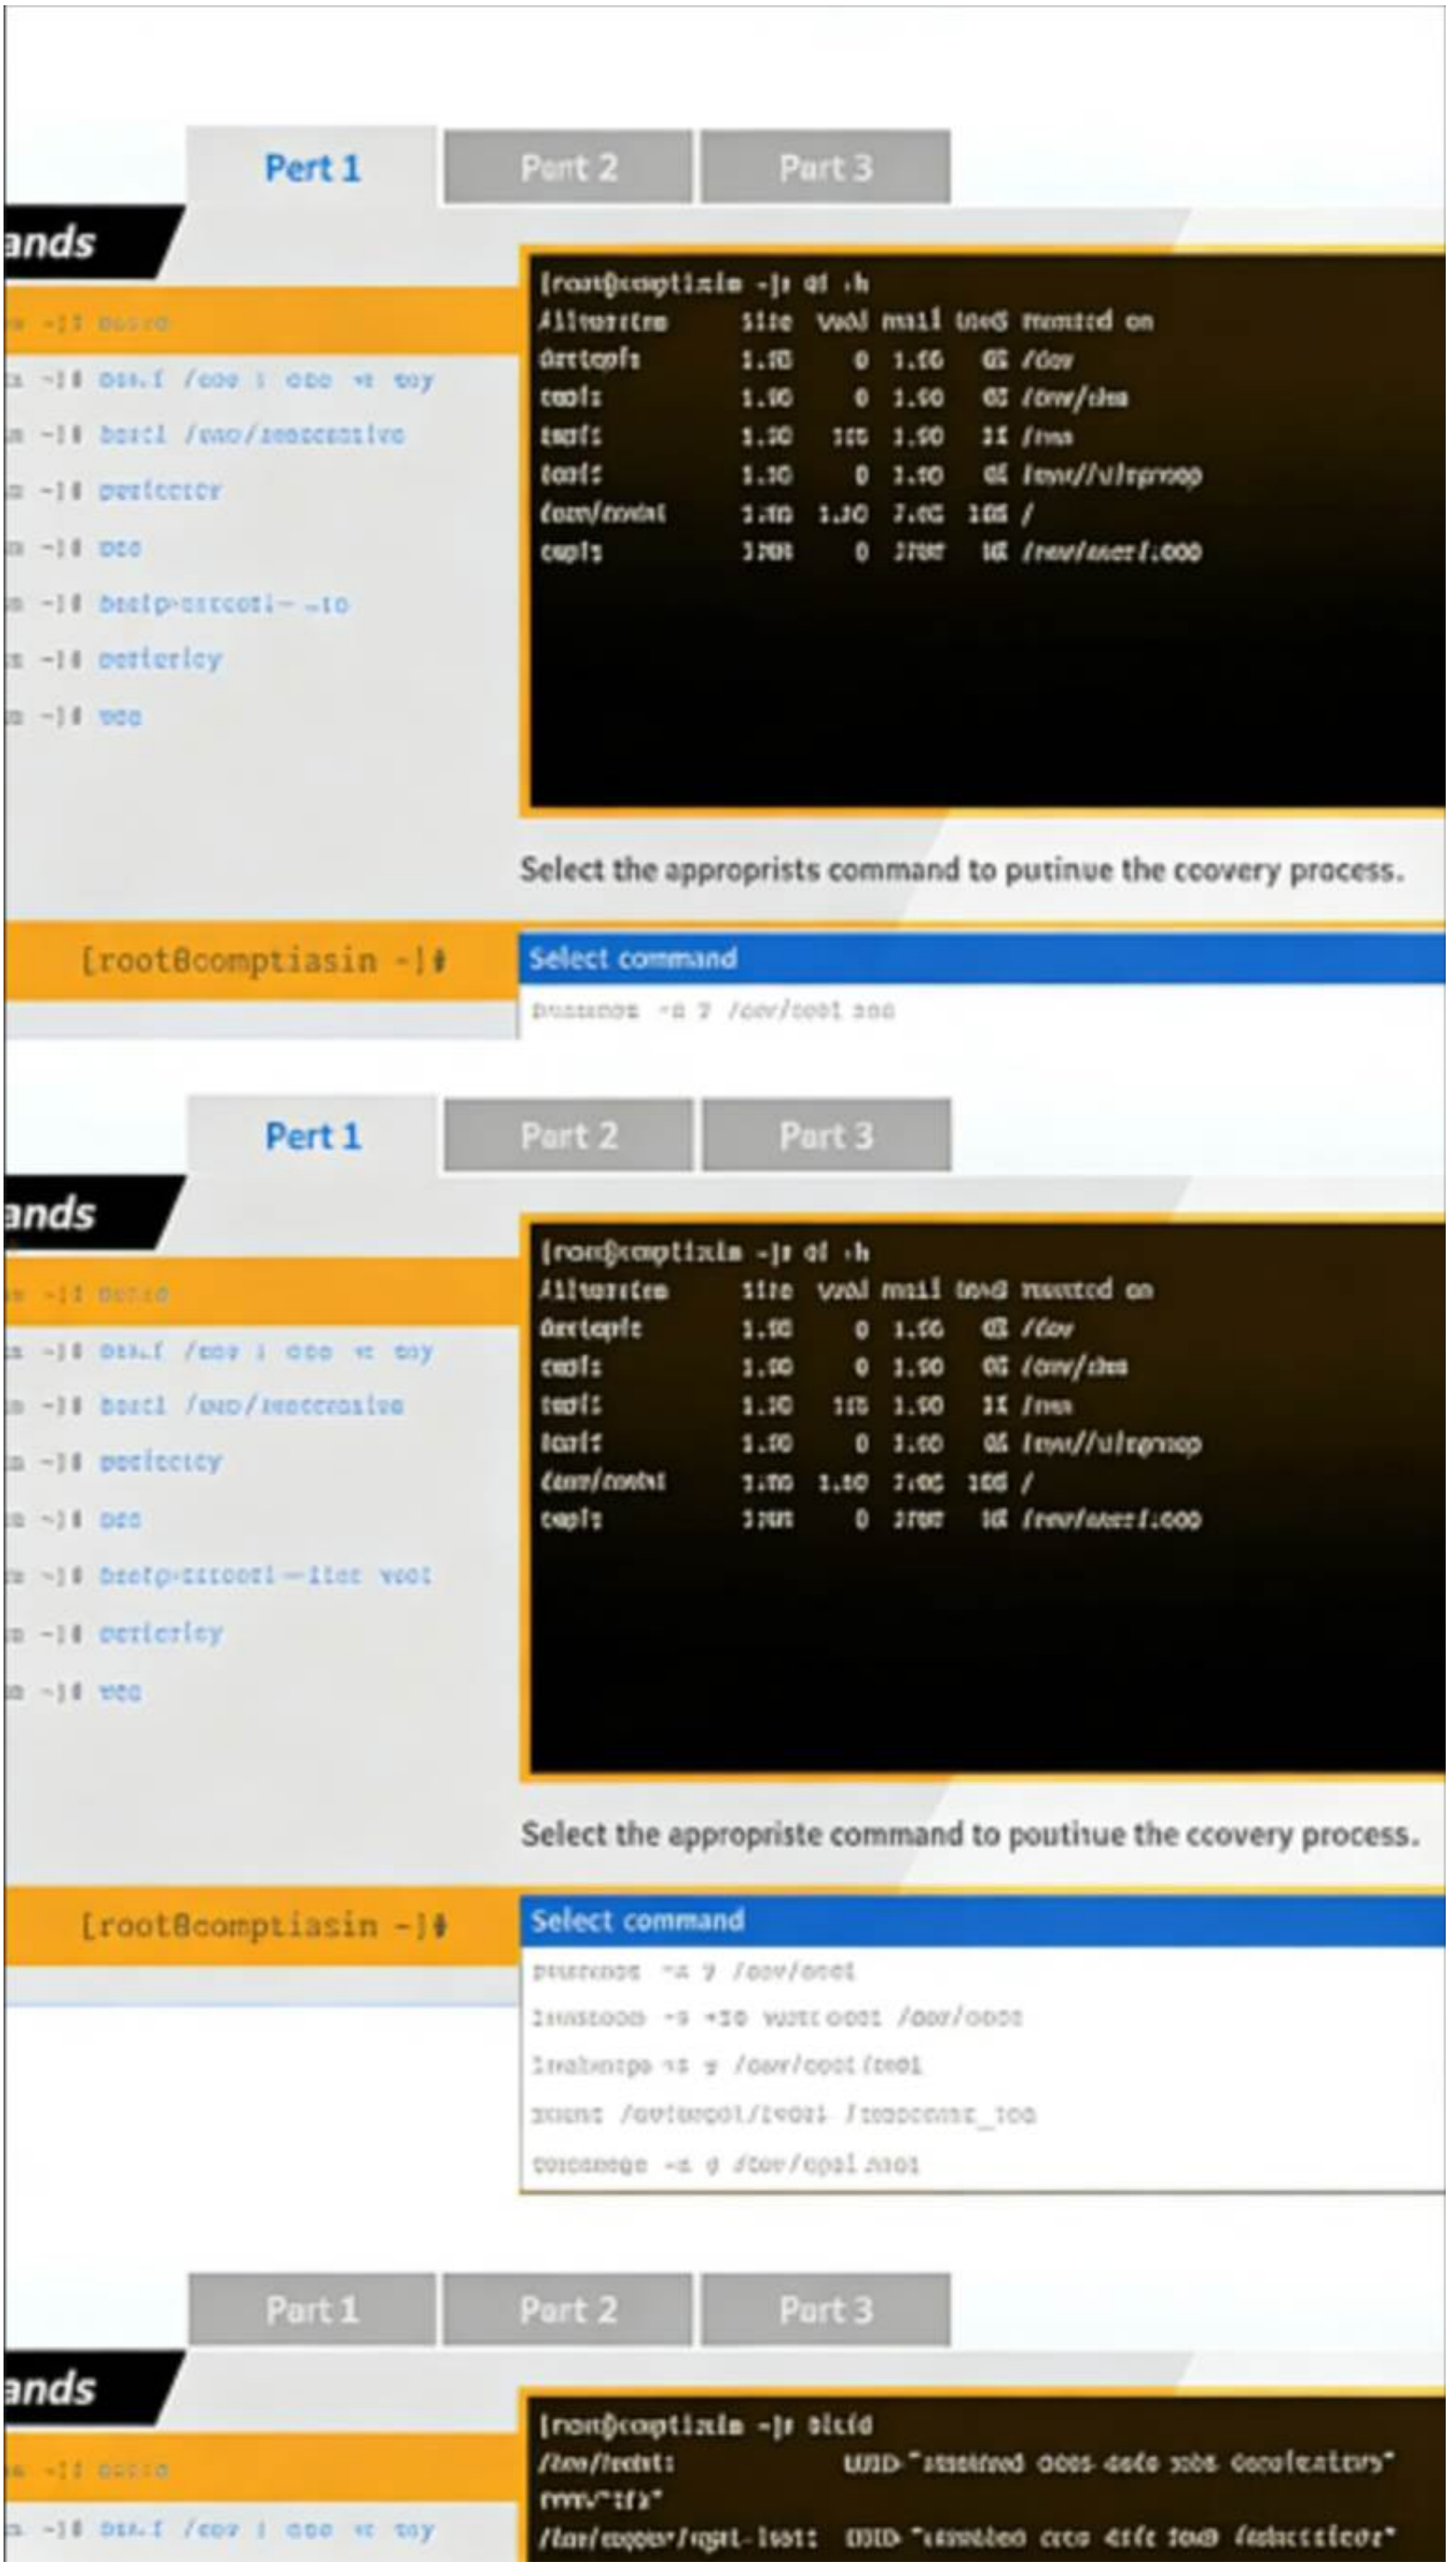
Part 1
Part 2
Part 3
Select the appropriate command to continue the recovery process.
[root@comptiasin ~]#
Select command
Part 1
Part 2
Part 3
Select the appropriate command to continue the recovery process.
[root@comptiasin ~]#
Select command
Part 1
Part 2
Part 3

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 (pvdisplay, 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 50
A Linux user needs to download the latest Debian image from a Docker repository. Which of the following commands makes this task possible?

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.

NEW QUESTION 55
A DevOps engineer made some changes to files in a local repository. The engineer realizes that the changes broke the application and the changes need to be reverted back. Which of the following commands is the best way to accomplish this task?

A. git pull
B. git reset
C. git rebase
D. git stash

Answer: B

Explanation:
Version control rollback operations are a core DevOps skill covered in the Linux+ V8 objectives. When changes in a local Git repository break an application and must be reverted, the administrator must choose a command that directly undoes those changes.
The command git reset is the most appropriate option in this scenario. It allows the engineer to move the current branch pointer (HEAD) to a previous commit, effectively discarding or undoing local changes. Depending on the reset mode (--soft, --mixed, or --hard), the engineer can control whether changes are preserved in the staging area or working directory. This flexibility makes git reset the primary tool for reverting problematic local changes.

The other options are not suitable. git pull fetches and merges changes from a remote repository and does not revert local modifications. git rebase rewrites commit history and is used to reapply commits on top of another base, not to undo broken changes. git stash temporarily saves uncommitted changes for later use but does not revert the repository to a stable state.
Linux+ V8 documentation emphasizes that git reset is commonly used during local development when changes need to be undone quickly before being shared with others. Therefore, the correct answer is B.

NEW QUESTION 57
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?

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.

NEW QUESTION 59
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 unconf ined_u context.
B. The administrator forgot to restart the SSHD after creating the authorizedjceys file.
C. The authorized_keys file has the incorrect root permissions assigned.
D. The authorized_keys file does not have the correct security context to match SELinux policy.

Answer: D

Explanation:
This issue is directly related to SELinux enforcement, which is a key topic in the Security domain of CompTIA Linux+ V8. The logs clearly indicate that SSH key-based authentication is failing due to an SELinux access control violation rather than a traditional file permission or SSH configuration problem.
The most important clue is the AVC denial message, which shows that the sshd process is being denied read access to the authorized_keys file. The security context of the file is listed as unconfined_u:object_r: home_root_t:s0. Under a targeted SELinux policy, SSH is only permitted to read authorized_keys files that are labeled with the correct SELinux type, typically ssh_home_t.
Because SELinux is running in enforcing mode, it actively blocks access that violates policy rules, even if standard UNIX permissions are correct. Although the file permissions (600) are acceptable for an authorized_keys file, SELinux does not rely solely on traditional permissions. The mismatch between the expected SELinux context and the actual context prevents sshd from accessing the file, causing SSH to fall back to password authentication.
Option D correctly identifies the root cause: the authorized_keys file does not have the correct SELinux security context. This is a well-documented Linux+ V8

troubleshooting scenario, commonly resolved by restoring the correct context using commands such as restorecon or by ensuring the file resides in a properly labeled home directory.
The other options are incorrect. Restarting sshd does not fix SELinux labeling issues. The policy itself is functioning as intended, and file ownership alone does not override SELinux access controls.
Linux+ V8 documentation emphasizes that SELinux denials must be addressed by correcting file contexts rather than weakening security controls. Therefore, the correct answer is D.

NEW QUESTION 62
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 67
A Linux administrator attempts to log in to a server over SSH as root and receives the following error message: Permission denied, please try again. The administrator is able to log in to the console of the server directly with root and confirms the password is correct. The administrator reviews the configuration of the SSH service and gets the following output:

```
Port 22
PermitRootLogin prohibit-password
PasswordAuthentication yes
PermitEmptyPassword no
Use PAM no
MaxSessions 1
MaxAuthTries 3
```

Based on the above output, which of the following will most likely allow the administrator to log in over SSH to the server?

A. Log out other user sessions because only one is allowed at a time.
B. Enable PAM and configure the SSH module.
C. Modify the SSH port to use 2222.
D. Use a key to log in as root over SSH.

Answer: D

Explanation:
The SSH configuration option PermitRootLogin prohibit-password prevents the root user from logging in with password authentication. This setting means root cannot use a password to log in via SSH; only key- based authentication is permitted for root. The administrator can still log in as root locally, which is not affected by this SSH configuration. To allow SSH access as root, the administrator must use an SSH key instead of a password.
Other options:
* A. MaxSessions controls the number of simultaneous SSH sessions but is not causing the login denial here.
* B. PAM (Pluggable Authentication Modules) is disabled, but enabling it is not required for basic SSH authentication.
* C. Changing the SSH port is unrelated to the authentication method issue.
Reference:
CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 11: "Securing Linux", Section: "Securing SSH Access"
CompTIA Linux+ XK0-006 Objectives, Domain 3.0: Security

NEW QUESTION 71
Which of the following describes PEP 8?

A. The style guide for Python code
B. Python virtual environments
C. A package installer for Python
D. A Python variable holding octal values

Answer: A

Explanation:
Python scripting is part of Linux automation, and Linux+ V8 includes knowledge of Python development standards. PEP 8 stands for Python Enhancement Proposal 8 and defines the official style guide for Python code.
PEP 8 provides conventions for code layout, indentation, naming, line length, whitespace usage, and commenting. Its purpose is to improve code readability and maintainability, especially in collaborative environments. Linux+ V8 emphasizes that standardized coding practices are critical in automation and DevOps workflows.
The other options are incorrect. Python virtual environments are managed using tools such as venv. Package installation is handled by pip. Octal values are represented using specific syntax and are unrelated to PEP 8.
Therefore, the correct answer is A.

NEW QUESTION 75
Which of the following passwords is the most complex?

A. H3sa1dt01d
B. he$@ID$heTold
C. H3s@1dSh3t0|d
D. HeSaidShetold

Answer: C

Explanation:
Comprehensive and Detailed 250 to 350 words of Explanation From Linux+ V8 documents:
Password complexity is a fundamental concept within theSecuritydomain of CompTIA Linux+ V8. Complex passwords significantly reduce the risk of successful brute-force, dictionary, and credential-stuffing attacks. Linux+ emphasizes evaluating passwords based onlength, character variety, unpredictability, and resistance to common word patterns.
OptionC, H3s@1dSh3t0|d, is the most complex password among the choices. It demonstrates strong security characteristics by incorporating:
Uppercase letters (H, S)
Lowercase letters (s, d, t)
Numbers (3, 1, 0)
Multiple special characters (@, |)
A longer overall length compared to some other options
Additionally, optionCusescharacter substitution (leet-style)in a way that breaks up recognizable words more effectively than the other choices. This significantly increases entropy and makes the password harder to guess using rule-based or hybrid cracking techniques.
OptionAincludes uppercase letters and numbers but lacks special characters and is relatively short. OptionBincludes special characters and mixed case, but it still closely resembles readable words, making it more susceptible to dictionary-based attacks. OptionDuses only alphabetic characters and clear word patterns, making it the weakest choice.
Linux+ V8 documentation highlights that thestrongest passwords combine length with diverse character classes and minimal predictability. PasswordCbest meets all of these criteria and would score highest against common password-cracking strategies.
Therefore, the correct answer isC. H3s@1dSh3t0|d.

NEW QUESTION 77
Which of the following can reduce the attack surface area in relation to Linux hardening?

A. Customizing the log-in banner
B. Reducing the number of directories created
C. Extending the SSH startup timeout period
D. Enforcing password strength and complexity

Answer: D

Explanation:
Comprehensive and Detailed Explanation From Exact Extract:
Reducing the attack surface area in Linux hardening refers to limiting possible points of unauthorized access. According to the CompTIA Linux+ Official Study Guide (Exam XK0-006), enforcing strong password policies is a critical aspect of security hardening. This practice ensures that user accounts are protected by passwords that are difficult to guess or crack, thus minimizing the risk of successful brute-force attacks. Implementing password complexity requirements (such as minimum length, use of uppercase, lowercase, numbers, and special characters) directly addresses one of the primary vectors for unauthorized access.
Other options do not have a direct impact on reducing the attack surface:
* A. Customizing the log-in bannerserves as a legal notification and does not affect system vulnerabilities.
* B. Reducing the number of directories createdis not related to hardening or access control.
* C. Extending the SSH startup timeout periodmay give attackers more time to attempt a connection and does not increase security.
[Reference:, CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 11: "Securing the System", Section: "Implementing Password Policies", CompTIA Linux+ XK0-006 Exam Objectives, Domain 3.0: Security, , ,]

NEW QUESTION 81
A Linux systems administrator needs to extract the contents of a file named /home/dev/web.bkp to the /var/www/html/ directory. Which of the following commands should the administrator use?

A. cd /var/www/html/ && gzip -c /home/dev/web.bkp | tar xf -
B. pushd /var/www/html/ && cpio -idv < /home/dev/web.bkp && popd
C. tar -c -f /home/dev/web.bkp /var/www/html/
D. unzip -c /home/dev/web.bkp /var/www/html/

Answer: B

Explanation:
Comprehensive and Detailed 250 to 350 words of Explanation From Linux+ V8 documents:
File extraction and backup restoration are fundamentalSystem Managementtasks covered in CompTIA Linux+ V8. In this scenario, the administrator must extract the contents of an existing backup file into a target directory.
The correct command isoption B, which uses cpio in extract mode. The command changes into the destination directory (/var/www/html/) using pushd, extracts the archive contents with cpio -idv, and then returns to the original directory with popd. This ensures that files are restored into the correct location without modifying

paths inside the archive.
The cpio utility is commonly used for backups created with cpio -o and supports reading archive data from standard input. Linux+ V8 documentation includes cpio as a valid and supported archive format for backup and restore operations.
The other options are incorrect. OptionAincorrectly assumes the backup is a gzip-compressed tar archive. OptionCcreates a new archive instead of extracting one. OptionDassumes the file is a ZIP archive, which is not indicated by the .bkp extension.
Linux+ V8 emphasizes using the correct tool based on the archive format and restoring files into the intended directory. Therefore, the correct answer isB.

NEW QUESTION 85
A systems administrator needs to set the IP address of a new DNS server. Which of the following files should the administrator modify to complete this task?

A. /etc/whois.conf
B. /etc/resolv.conf
C. /etc/nsswitch.conf
D. /etc/dnsmasq.conf

Answer: B

Explanation:
DNS client configuration is a foundational Linux networking task covered in Linux+ V8 system management objectives. When an administrator needs to specify the IP address of a DNS server that the system should use for name resolution, the correct file to modify is/etc/resolv.conf.
The /etc/resolv.conf file defines DNS resolver settings, including one or more nameserver entries that specify the IP addresses of DNS servers. Applications and system services rely on this file to resolve hostnames to IP addresses.
The other options are incorrect. /etc/whois.conf configures WHOIS queries. /etc/nsswitch.conf controls the order of name resolution sources but does not define DNS server IP addresses. /etc/dnsmasq.conf configures a local DNS caching service, not the system-wide resolver directly.
Linux+ V8 documentation highlights /etc/resolv.conf as the authoritative DNS client configuration file, though it may be dynamically managed by tools such as NetworkManager or systemd-resolved.
Therefore, the correct answer isB. /etc/resolv.conf.

NEW QUESTION 86
To perform a live migration, which of the following must match on both host servers?(Choose two)

A. USB ports
B. Network speed
C. Available swap
D. CPU architecture
E. Available memory
F. Disk storage path

Answer: DE

Explanation:
Comprehensive and Detailed 250 to 350 words of Explanation From Linux+ V8 documents:
Live migration is a virtualization feature that allows a running virtual machine to be moved from one host to another with minimal or no downtime. This topic falls underSystem Managementin the CompTIA Linux+ V8 objectives, particularly in the areas of virtualization and resource management.
For a live migration to succeed, theCPU architecture must matchbetween the source and destination hosts. This is critical because the running virtual machine??s CPU state, instruction set, and registers must be compatible with the destination system. Migrating between different CPU architectures (for example, x86_64 to ARM) is not supported and would cause the virtual machine to fail. Therefore, optionDis required.
Additionally, the destination host must havesufficient available memoryto accommodate the virtual machine being migrated. During live migration, the memory contents of the running VM are copied from the source host to the destination host while the VM continues to run. If enough memory is not available, the migration cannot complete successfully. This makes optionEmandatory.
The other options are not strict requirements. USB ports do not need to match for live migration. Network speed may affect migration performance but does not need to be identical. Available swap space is not directly required for migration. Disk storage paths do not need to match as long as shared storage or compatible storage access is available.
Linux+ V8 documentation emphasizes CPU compatibility and memory availability as core prerequisites for live migration. Therefore, the correct answers areD and E.

NEW QUESTION 90
An administrator receives the following output while attempting to unmount a filesystem:
umount /data1: target is busy.
Which of the following commands should the administrator run next to determine why the filesystem is busy?

A. ps -f /data1
B. du -sh /data1
C. top -d /data1
D. lsof | grep /data1

Answer: D

Explanation:
Filesystem unmount failures are common troubleshooting scenarios covered in Linux+ V8. When the error"target is busy"appears, it means one or more processes are actively using files or directories within the mount point.
The correct diagnostic command islsof | grep /data1. The lsof (list open files) utility displays all open files and the processes using them. Filtering the output with grep /data1 identifies exactly which processes are holding file descriptors on the filesystem, preventing it from being unmounted.
The other options are incorrect. ps -f displays process information but does not show open file usage. du -sh calculates disk usage and does not identify active processes. top monitors system performance but cannot pinpoint filesystem locks.
Linux+ V8 documentation emphasizes using lsof or fuser to identify resource locks before unmounting filesystems. Therefore, the correct answer isD.

NEW QUESTION 91
A systems administrator needs to check the statuses of all the services on a Linux server. Which of the following commands accomplishes this task?

A. systemctl is-active --services
B. systemctl list-sockets --type=services
C. systemctl is-enabled --services
D. systemctl list-units --type=services

Answer: D

Explanation:
Service management usingsystemdis a core Linux+ V8 system management objective. Administrators frequently need to view the current status of all services to determine which ones are running, stopped, failed, or inactive.
The correct command issystemctl list-units --type=services, which displays all loaded service units along with their current state, including whether they are active, inactive, failed, or running. This provides a comprehensive, real-time view of service statuses on the system and is commonly used during troubleshooting and audits.
OptionA, systemctl is-active, is designed to check the status of a single service, not all services. OptionBlists socket units, not services. OptionC, systemctl is-enabled, checks whether services are enabled at boot, not whether they are currently running.
Linux+ V8 documentation explicitly references systemctl list-units --type=service as the primary command for viewing service runtime states. Therefore, the correct answer isD.

NEW QUESTION 94
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 97
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 emphasizessecurity, compliance, and governancewhen 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.
OptionBis 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 isB.

NEW QUESTION 102
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 involvesshared 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 togroup ownership inheritance.
By default, when a user creates a file or directory, it is owned by the user and assigned the user??sprimary 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 toset the setgid (set group ID) bit on the shared directory, making optionDcorrect. When the setgid bit is applied to a directory, all newly created files and subdirectories inherit thegroup 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. OptionA(setuid) is intended for executables, not directories. OptionBrequires constant manual intervention and does not scale. OptionCweakens security by granting write access to all users, violating the principle of least privilege.
Linux+ V8 documentation explicitly recommends using thesetgid bit on shared directoriesto manage collaborative access securely and efficiently.

NEW QUESTION 104
An administrator wants to search a file named myFile and look for all occurrences of strings containingat least five characters, wherecharacters two and five are i, butcharacter three is not b. Which of the following commands should the administrator execute to get the intended result?

A. grep .a*^b-.a myFile
B. grep .a., [a] myFile
C. grepa^b*a myFile
D. grep .i[^b].i myFile

Answer: D

Explanation:
Pattern matching using regular expressions is a key troubleshooting and text-processing skill covered in CompTIA Linux+ V8. The grep command, combined with regular expressions, allows administrators to search for complex string patterns within files.
The requirement specifies:
The string must containat least five characters
Character 2must be i
Character 3mustnotbe b
Character 5must 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
OptionD, 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 isD.

NEW QUESTION 108
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 109
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 112
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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