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

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

Answer: A

Explanation:
Comprehensive and Detailed 250 to 350 words of Explanation: From Linux+ V8 documents:
Container lifecycle management is a core topic within the Automation, Orchestration, and Scripting domain of CompTIA Linux+ V8. Administrators must understand the difference between creating containers, starting containers, and executing commands within running containers.
The correct command is docker run -it app-01-image. The docker run command performs three actions at once: it creates a new container from the specified image, starts the container, and optionally attaches the administrator??s terminal to it. The -i option keeps standard input open, while the -t option allocates a pseudo- terminal (TTY). Together, these options allow the administrator to interactively connect to the container immediately after it is created.
The other options are incorrect for the following reasons. docker start is used only to start an existing stopped container and does not create a new container from an image. Additionally, -t and -d are not valid options for attaching an interactive terminal during container startup. docker build is used to build a Docker image from a Dockerfile and cannot be used to create or connect to a container. docker exec is used to run commands inside an already running container and therefore cannot be used to create a container.

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

NEW QUESTION 5
A systems administrator is creating a backup copy of the /home/ directory. Which of the following commands allows the administrator to archive and compress the directory at the same time?

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

Answer: C

Explanation:
Creating backups is a core responsibility in Linux system management, and the Linux+ V8 objectives emphasize proper use of archiving and compression tools. The tar utility is the standard Linux tool for creating archive files, and it also supports compression through various options.
The command tar -cJf /backups/home.tar.xz /home/ correctly combines both archiving and compression in a single step. The -c option creates a new archive, -J specifies XZ compression, and -f allows the administrator to define the output file name. This results in a compressed archive of the entire /home/ directory, which is efficient for storage and transfer.
The other options are incorrect. cpio is an archiving tool but does not perform compression by itself without additional commands or pipelines. rsync -z compresses data during transfer but does not create an archive file. The dd command performs low-level copying of raw data and is not suitable for directory-based backups.
Linux+ V8 documentation highlights tar as the preferred utility for filesystem backups due to its flexibility, reliability, and support for multiple compression algorithms. Therefore, the correct answer is C.

NEW QUESTION 6
Which of the following most accurately describes a webhook?

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

Answer: D

Explanation:
Webhooks are commonly used in automation and DevOps workflows, which are emphasized in the Linux+ V8 objectives. A webhook is best described as 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 7
In the echo "profile-$num-$name" line of a shell script, the variable $num seems to not be expanding during execution. Which of the following notations ensures the value is expanded?

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

Answer: D

Explanation:
Shell variable expansion is a fundamental scripting concept included in Linux+ V8 objectives. In Bash and similar shells, variables are expanded only when they are interpreted within double quotes or unquoted contexts, and sometimes explicit syntax is required to avoid ambiguity.
The correct notation is ${num}, as shown in option D. Using curly braces around the variable name ensures the shell correctly identifies the variable boundary, especially when it is adjacent to other characters. This guarantees proper expansion of the variable's value.
The other options are incorrect. Single quotes prevent variable expansion entirely. The $(...) syntax is used for command substitution, not variable expansion. Quoting the variable name itself also prevents expansion.
Linux+ V8 documentation emphasizes ${VAR} notation as a best practice in shell scripting for clarity and correctness. Therefore, the correct answer is D.

NEW QUESTION 8
Which of the following commands should a Linux administrator use to determine the version of a kernel module?

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 itsversion, author, description, license, filename, and dependencies—the correct tool is modinfo.
The commandmodinfo bluetoothdisplays 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 isD. modinfo bluetooth.

NEW QUESTION 9
Which of the following utilities supports the automation of security compliance and vulnerability management?

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

Answer: D

Explanation:
Security compliance and vulnerability management are critical components of Linux system administration, and CompTIA Linux+ V8 places strong emphasis on automated security assessment tools.OpenSCAPis specifically designed to address these requirements.
OpenSCAP is an open-source framework that implements theSecurity Content Automation Protocol (SCAP), a set of standards used for automated vulnerability scanning, configuration compliance checking, and security auditing. It allows administrators to assess Linux systems against established security baselines such asCIS benchmarks, DISA STIGs, and organizational security policies. This makes OpenSCAP the most appropriate tool for automating both compliance and vulnerability management.
The other options serve different security-related purposes but do not fulfill the automation requirement.SELinuxis a mandatory access control system that enforces security policies at runtime but does not perform compliance scanning or vulnerability assessments.Nmapis a network scanning and discovery tool used to identify open ports and services, not compliance automation.AIDE(Advanced Intrusion Detection Environment) is a file integrity monitoring tool that detects unauthorized file changes but does not evaluate overall system compliance.
Linux+ V8 documentation highlights OpenSCAP as a tool used to automate security audits, generate compliance reports, and integrate with configuration management workflows. Its ability to standardize security checks across multiple systems makes it essential in enterprise and regulated environments.
Therefore, the correct answer isD. OpenSCAP.

NEW QUESTION 10
An administrator needs to verify the user ID, home directory, and assigned shell for the user named "accounting." Which of the following commands should the administrator use to retrieve this information?

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

Answer: A

Explanation:
User account information is centrally stored in the system??s account databases, and Linux+ V8 emphasizes the use of standard tools to query this data safely and consistently.
The getent passwd accounting command retrieves the user??s entry from 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 10
A systems administrator wants to review the amount of time the NetworkManager service took to start. Which of the following commands accomplishes this goal?

A. resolvectl
B. journalctl
C. systemctl daemon-reload
D. systemd-analyze blame

Answer: D

Explanation:
System boot performance analysis is an important system management task included in Linux+ V8. When administrators need to determine how long services take to start during boot,systemd analysis toolsare required.
The correct command issystemd-analyze blame. This command lists all systemd services and shows how long each one took to initialize during the boot process. It is commonly used to identify slow-starting services that may impact system startup performance, including NetworkManager.
The other options are incorrect. resolvectl is used for DNS resolution management and provides no service timing information. journalctl can display logs but does not provide a clear, summarized service startup timing report. systemctl daemon-reload only reloads systemd unit files and does not perform analysis.
Linux+ V8 documentation explicitly references systemd-analyze blame as the correct tool for diagnosing service startup delays. Therefore, the correct answer isD.

NEW QUESTION 15
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 16
A Linux administrator tries to install Ansible in a Linux environment. One of the steps is to change the owner and the group of the directory /opt/Ansible and its contents. Which of the following commands will accomplish this task?

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

Answer: B

Explanation:
Comprehensive and Detailed Explanation From Exact Extract:
The chown command is used to change the owner and group of files and directories. The -R (recursive) flag ensures that all contents within the directory are also updated. The correct syntax is chown -R owner:group directory. So, chown -R Ansible:Ansible /opt/Ansible will change the owner and group for /opt/Ansible and everything inside it to "Ansible".
Other options:
* A.groupmod is used to modify group properties, not ownership of directories or files.
* C.usermod is for modifying user properties or group memberships.
* D.chmod changes permissions, not owner/group.
[Reference:, CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 6: "User and Group Management", Section: "Managing File Ownership and Permissions", CompTIA Linux+ XK0-006 Objectives, Domain 1.0: System Management,]

NEW QUESTION 21
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 22
A new drive was recently added to a Linux system. Using the environment and tokens provided, complete the following tasks:
• Create an appropriate device label.
• Format and create an ext4 file system on the new partition. The current working directory is /.

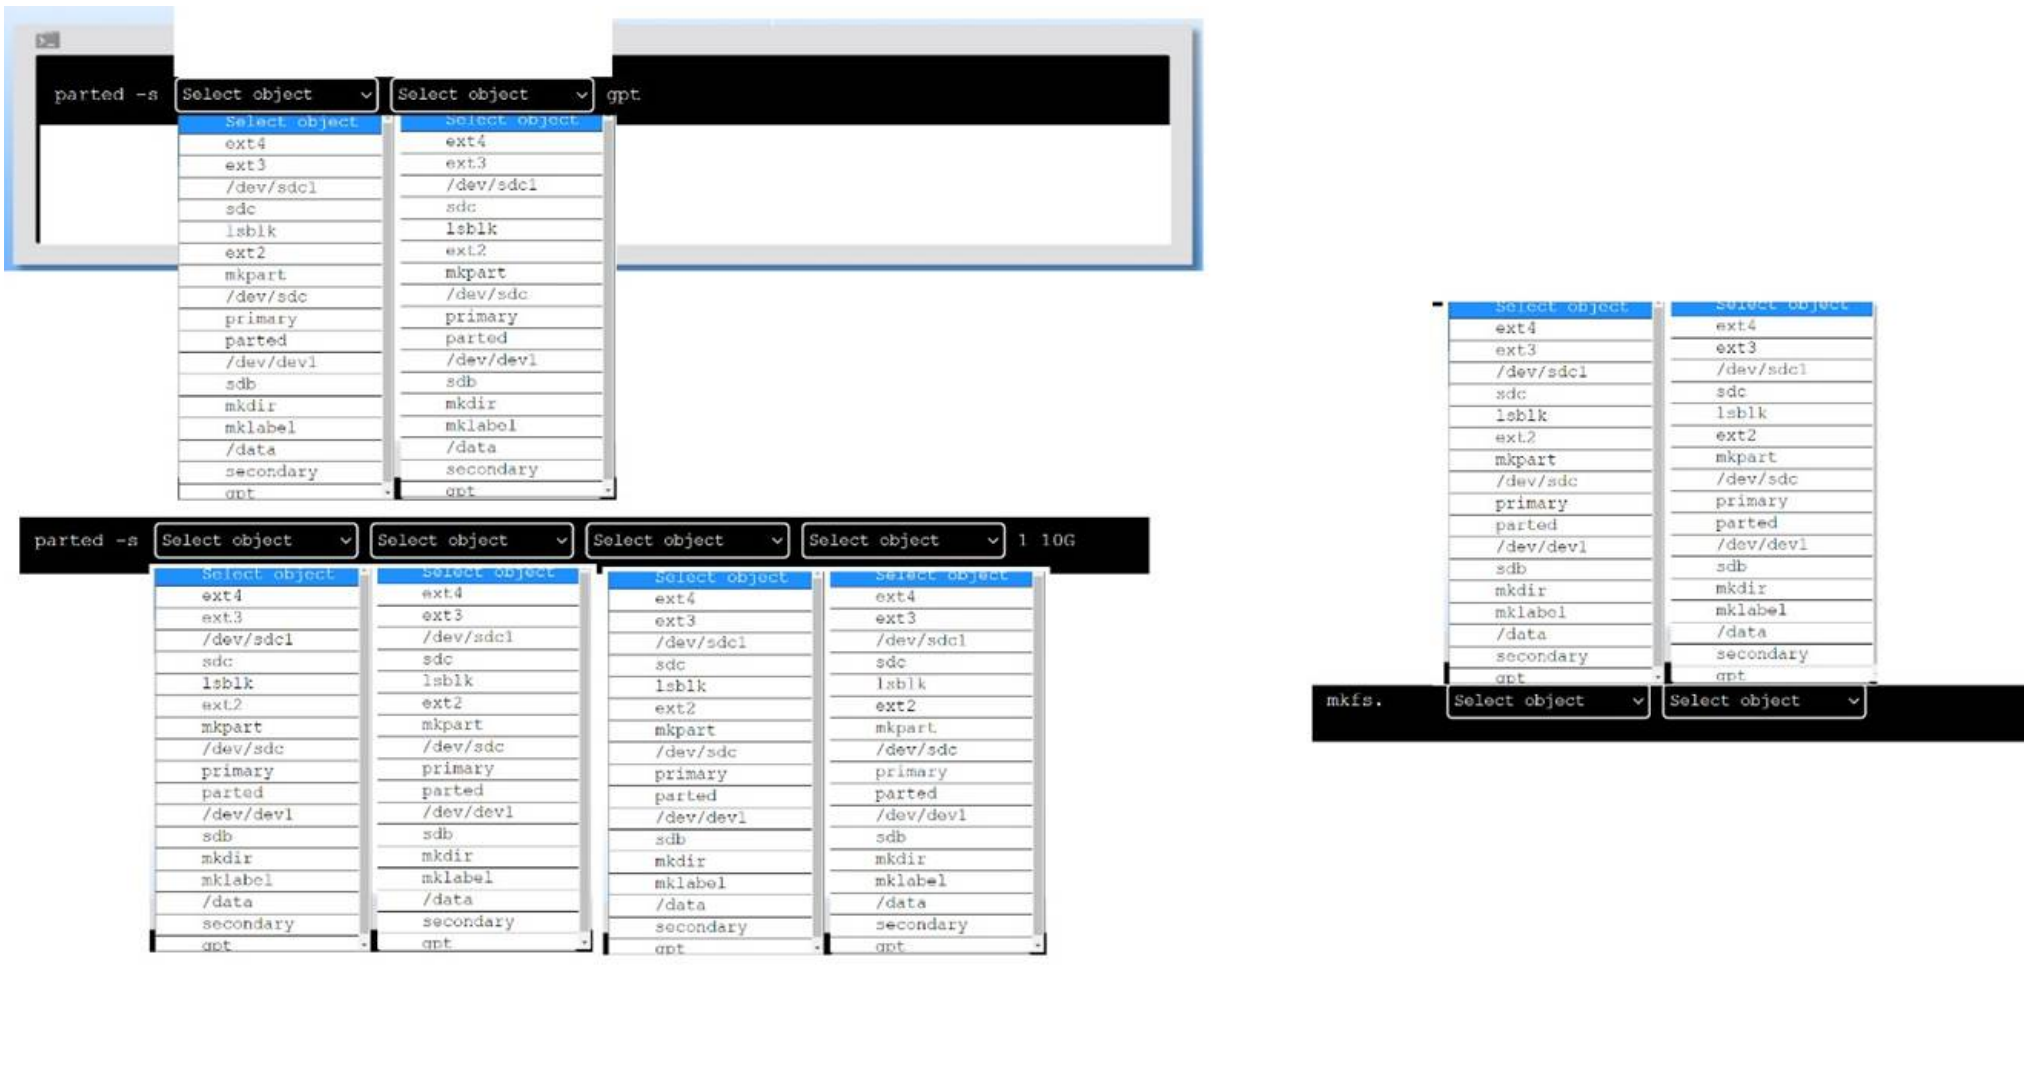

A. Mastered
B. Not Mastered

Answer: A

Explanation:
To create an appropriate device label, format and create an ext4 file system on the new partition, you can use the following commands:
To create a GPT (GUID Partition Table) label on the new drive /dev/sdc, you can use the parted command with the -s option (for script mode), the device name (/dev/sdc), the mklabel command, and the label type (gpt). The command is:
parted -s /dev/sdc mklabel gpt
To create a primary partition of 10 GB on the new drive /dev/sdc, you can use the parted command with the -s option, the device name (/dev/sdc), the mkpart command, the partition type (primary), the file system type (ext4), and the start and end points of the partition (1 and 10G). The command is:
parted -s /dev/sdc mkpart primary ext4 1 10G
To format and create an ext4 file system on the new partition /dev/sdc1, you can use the mkfs command with the file system type (ext4) and the device name (/dev/sdc1). The command is:
mkfs.ext4 /dev/sdc1
You can verify that the new partition and file system have been created by using the lsblk command, which will list all block devices and their properties.

NEW QUESTION 25
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 30
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 32
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 33
An administrator needs to remove the directory /home/user1/data and all of its contents. Which of the following commands should the administrator use?

A. rmdir -p /home/user1/data
B. ln -d /home/user1/data
C. rm -r /home/user1/data
D. cut -d /home/user1/data

Answer: C

Explanation:
File and directory management is a core system administration skill addressed in Linux+ V8. When an administrator needs to delete a directory that contains files or subdirectories, a recursive deletion is required.
The correct command is rm -r /home/user1/data. The rm command removes files, and the -r (recursive) option allows it to delete directories and all of their contents, including nested files and subdirectories. This is the standard and correct method for removing non-empty directories.
The other options are incorrect. rmdir -p only removes empty directories and will fail if the directory contains files. ln -d is used to create directory hard links, not remove directories. cut -d is a text-processing command unrelated to filesystem operations.
Linux+ V8 documentation stresses caution when using rm -r, as it permanently deletes data without recovery unless backups exist. Therefore, the correct answer

is C.

NEW QUESTION 38
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 39
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 40
A Linux user frequently tests shell scripts located in the /home/user/scripts directory. Which of the following commands allows the user to run the program by invoking only the script name?

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

Answer: C

Explanation:
In Linux, the ability to execute a program by typing only its name depends on whether the directory containing the executable is included in the user??s PATH environment variable. The PATH variable defines a colon-separated list of directories that the shell searches when a command is entered.
Option C, export PATH=$PATH:/home/user/scripts, correctly appends the /home/user/scripts directory to the existing PATH variable. Once this command is executed, any executable script located in that directory can be run simply by typing its filename, provided the script has execute permissions. This behavior is explicitly covered in the Linux+ V8 objectives related to environment variables and shell configuration.
The other options are incorrect. Option A incorrectly attempts to redefine the SHELL variable and uses invalid syntax. Option B modifies the TERM variable, which controls terminal type and has nothing to do with command execution. Option D attempts to create an alias using invalid syntax and would not affect command lookup behavior.
Linux+ V8 documentation emphasizes modifying the PATH variable as the standard and recommended method for simplifying script execution. This approach is commonly used by developers and administrators who frequently run custom scripts.
Therefore, the correct answer is C.

NEW QUESTION 41
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 45
Which of the following describes how a user's public key is used during SSH authentication?

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 key
C. If it is found, the user is allowed to log in.
D. The user's public key is used instead of a password to allow server access.
E. 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.
Reference:
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
==========

NEW QUESTION 47
Users cannot access a server after it has been restarted. At the server console, the administrator runs the following commands;

```
$ ss -lnt
 State      Recv-Send-Q     Q          LocalAddress:PortPeerAddress:PortProcess
LISTEN     0    32     0.0.0.0:53    0.0.0.0:*
LISTEN     0   128     0.0.0.0:22    0.0.0.0:*
LISTEN     0  1024    0.0.0.0:443    0.0.0.0:*
LISTEN     0  4096   0.0.0.0:5355    0.0.0.0:*
LISTEN     0     5127.0.0.1:4711    0.0.0.0:*

$ sudo firewall-cmd --list-all
FedoraServer (active)
  target: default
  icmp-block-inversion: no
  interfaces: enp3s0
  sources:
  services: cockpit dhcp dhcpv6-client dns dns-over-tls https
  [...]

$ uptime
14:52:35 up 1 day, 3:08, 1 user, load average: 0.05, 0.07, 0.07

$ ping server1 -c 5
PING server1 (192.168.0.2) 56(84) bytes of data.
64 bytes from server1 (192.168.0.2): icmp_seq=1 ttl=64 time=0.436 ms
64 bytes from server1 (192.168.0.2): icmp_seq=2 ttl=64 time=0.644 ms
...
```

Which of the following is the cause of the issue?

A. The DNS entry does not have a valid IP address.
B. The SSH service has not been allowed on the firewall.
C. The server load average is too high.
D. The wrong protocol is being used to connect to the web server.

Answer: B

Explanation:
This issue is a classic example of post-reboot connectivity troubleshooting, which falls under the Troubleshooting domain of CompTIA Linux+ V8. The administrator has correctly gathered evidence using multiple diagnostic tools, allowing the root cause to be identified through correlation.
The ss -lnt output confirms that the SSH daemon is running and listening on TCP port 22. This eliminates the possibility that the SSH service failed to start after reboot. Additionally, the uptime output shows a very
low load average, indicating that system performance is not a limiting factor. The successful ping test confirms that the server is reachable at the network layer and that DNS resolution and basic connectivity are functioning correctly.
The critical clue comes from the firewall configuration. The output of firewall-cmd --list-all shows that only specific services are allowed through the firewall, such as https, dns, and cockpit. The SSH service is notably absent. On systems using firewalld, services must be explicitly allowed, even if the daemon itself is running and listening on the correct port.
As a result, incoming SSH connection attempts are being blocked by the firewall, preventing users from accessing the server remotely after reboot. This aligns precisely with option B.
The other options are incorrect. DNS is functioning, as shown by successful ping responses. System load is low and not contributing to the issue. There is no indication that users are attempting to access the web server using an incorrect protocol.
Linux+ V8 documentation emphasizes that administrators must verify both service status and firewall rules when diagnosing access issues. In this case, allowing SSH with a command such as firewall-cmd --add- service=ssh --permanent followed by a reload would resolve the problem.

NEW QUESTION 49
A Linux administrator is testing a web application on a laboratory service and needs to temporarily allow DNS and HTTP/HTTPS traffic from the internal network. Which of the following commands will accomplish this task?

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

Answer: C

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:
The correct way to temporarily allow specific services in a particular zone with firewalld is to use firewall- cmd --add-service=service --zone=zone. Multiple services can be specified in curly braces and separated by commas. The correct syntax is:
bash CopyEdit
firewall-cmd --add-service={dns,http,https} --zone=internal
This command will allow DNS (port 53), HTTP (port 80), and HTTPS (port 443) through the firewall for the "internal" zone temporarily (for the current runtime session).
Other options:
* A. The command syntax is incorrect; firewalld is a service, not a command-line tool.
* B. iptables does not use the --enable-service flag, nor does it have zones in this way.
* D. systemctl mask disables services, and the rest of the command is invalid.
Reference:
CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 9: "Networking", Section: "Managing Firewalls with firewalld"
CompTIA Linux+ XK0-006 Objectives, Domain 2.0: Networking
==========

NEW QUESTION 53
Which of the following Ansible components contains a list of hosts and host groups?

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

Answer: B

Explanation:
Ansible architecture and core components are part of the Automation, Orchestration, and Scripting domain in CompTIA Linux+ V8. Among these components, the inventory plays a foundational role in defining the infrastructure Ansible manages.
An Ansible inventory is a file (or set of files) that contains a list of managed hosts and optionally organizes them into logical groups. These hosts can be defined by IP address, fully qualified domain name (FQDN), or hostname. Inventories may be written in INI, YAML, or dynamically generated formats. Grouping hosts allows administrators to apply configurations, roles, and tasks to multiple systems simultaneously.
Option B, Inventory, is correct because it explicitly defines which systems Ansible will target. Without an inventory, Ansible does not know where to execute tasks. Linux+ V8 documentation emphasizes inventories as the starting point for all Ansible operations.
The other options are incorrect. Facts are system variables automatically collected by Ansible about managed hosts, such as OS version or IP address. Playbooks define what actions to perform but rely on the inventory to know where to perform them. Collections are distribution units that package roles, modules, and plugins, not host definitions.
Therefore, the correct answer is B. Inventory.

NEW QUESTION 58
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 63
Which of the following commands should an administrator use to convert a KVM disk file to a different format?

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

Answer: D

Explanation:
Virtualization management is part of Linux system administration and is included in Linux+ V8 objectives. KVM virtual machines commonly use disk image formats such as qcow2, raw, or vmdk. Converting between these formats is a routine administrative task.
The correct tool for disk image conversion isqemu-img. This utility allows administrators to create, convert, resize, and inspect virtual disk images. For example, converting a qcow2 image to raw format can be accomplished using qemu-img convert. This capability is explicitly referenced in Linux+ V8 documentation related to virtualization tooling.
The other options are incorrect. qemu-kvm refers to the hypervisor component, not disk manipulation. qemu-ng is not a valid QEMU utility. qemu-io is used for low-

level I/O testing and debugging, not image format conversion.
Therefore, the correct answer isD. qemu-img.

NEW QUESTION 64
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 68
Which of the following describes the method of consolidating system events to a single location?

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

Answer: A

Explanation:
Comprehensive and Detailed 250 to 350 words of Explanation From Linux+ V8 documents:
Consolidating system events from multiple sources into a single, centralized location is a key concept in Linux system administration and is explicitly covered underlogging and monitoringtopics in the CompTIA Linux+ V8 objectives. This method is known aslog aggregation, making optionAthe correct answer.
Log aggregation refers to the practice of collecting logs generated by operating systems, services, applications, and network devices and storing them in a centralized repository. In Linux environments, logs may originate from systemd-journald, syslog, application-specific log files, containers, and cloud-based workloads. Aggregating these logs allows administrators to analyze events more efficiently, correlate issues across systems, and improve troubleshooting, auditing, and security monitoring.
Linux+ V8 documentation emphasizes centralized logging as a best practice in environments with multiple servers. Without log aggregation, administrators would need to log in to each system individually to inspect logs, which is inefficient and error-prone. Centralized solutions such as syslog servers, ELK/EFK stacks, and SIEM platforms enable real-time analysis, long-term retention, and alerting based on log data.
The other options do not describe log consolidation.Health checksare used to verify whether services or systems are operational but do not collect or store event data.Webhooksare HTTP-based callbacks used for event-driven automation and notifications, not for storing logs.Threshold monitoringinvolves generating alerts when metrics exceed defined limits, such as CPU or memory usage, but it does not centralize system event records.
Linux+ V8 stresses that effective log aggregation improves incident response, supports compliance requirements, and enhances system visibility. It is especially important for detecting security incidents, diagnosing failures, and performing root-cause analysis across distributed systems.

NEW QUESTION 69
Following the completion of monthly server patching, a Linux administrator receives reports that a critical application is not functioning. Which of the following commands should help the administrator determine which packages were installed?

A. dnf history
B. dnf list
C. dnf info
D. dnf search

Answer: A

Explanation:
Package management troubleshooting is a critical Linux administration skill addressed in CompTIA Linux+ V8. After system patching, identifying which packages were installed, updated, or removed is often the first step in diagnosing application failures.
Thednf historycommand is specifically designed for this purpose. It displays a chronological list of all DNF transactions, including installations, upgrades, downgrades, and removals. Each transaction is assigned an ID and includes timestamps, affected packages, and actions taken. This allows administrators to correlate application failures with recent changes.
OptionAis correct because it provides historical context rather than just current package state. Linux+ V8 documentation highlights dnf history as an essential auditing and rollback tool.
The other options are insufficient. dnf list shows installed or available packages but does not indicate when they were installed. dnf info displays metadata for a specific package but does not show transaction history. dnf search is used to find packages by name or description.
By reviewing recent transactions with dnf history, administrators can quickly identify problematic updates and take corrective action, such as rolling back a package.
Therefore, the correct answer isA.

NEW QUESTION 73
A systems administrator receives reports about connection issues to a secure web server. Given the following firewall and web server outputs:
Firewall output:

Status: active
To Action From
443/tcp DENY Anywhere
443/tcp (v6) DENY Anywhere (v6)
Web server output:
tcp LISTEN 0 4096 *:443 :
Which of the following commands best resolves this issue?

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

Answer: C

Explanation:
This scenario involves firewall configuration and service accessibility, which falls under the Security domain of the CompTIA Linux+ V8 objectives. The key to resolving this issue is interpreting both the firewall output and the web server status correctly.
The web server output shows that the service is actively listening on TCP port 443, which is the standard port for HTTPS (secure web traffic). The line tcp LISTEN 0 4096 *:443 *:* confirms that the web server is running properly and is ready to accept incoming connections on port 443 from any interface. This indicates that the problem is not with the web server configuration itself.
However, the firewall output clearly shows that incoming connections to port 443 are being blocked. The rules 443/tcp DENY Anywhere and 443/tcp (v6) DENY Anywhere (v6) indicate that the Uncomplicated Firewall (UFW) is explicitly denying HTTPS traffic for both IPv4 and IPv6. As a result, external clients cannot establish a secure connection to the server, even though the service is running correctly.
To resolve this issue securely and correctly, the administrator must remove the firewall rule that denies HTTPS traffic. Option C, ufw delete deny https/tcp, directly removes the blocking rule while preserving the rest of the firewall configuration. This aligns with Linux+ best practices, which emphasize making precise firewall changes rather than disabling security controls entirely.
The other options are incorrect. Option A, ufw disable, would completely turn off the firewall, creating a significant security risk. Option B, ufw allow 80/tcp, only opens HTTP traffic on port 80 and does not resolve HTTPS connectivity issues. Option D, ufw allow 4096/tcp, incorrectly attempts to open an internal socket backlog value rather than a valid service port.
Therefore, the correct and most secure solution is C.

NEW QUESTION 75
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 78
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 81
A systems administrator wants to prevent the current contents of a file from being overwritten and wants to allow new additions at the end of the file. Which of the following commands should the administrator use?

A. setenforce file
B. setfacl -m m::t file
C. chattr +a file
D. chmod +t file

Answer: C

NEW QUESTION 84
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 86
An administrator set up a new user account called "test". However, the user is unable to change their password. Given the following output:

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

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

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

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

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

A. The SUID bit is missing on the /bin/passwd file.
B. The password provided by the user "test" does not meet complexity requirements.
C. The user "test" already changed the password today.
D. The password has been disabled for user "test".

Answer: A

Explanation:
For normal users to change their own password, /bin/passwd must have the SUID bit set (permissions should be -rwsr-xr-x). The SUID bit allows users to run the program with the permissions of the file owner (root), which is required to update /etc/shadow. The provided output shows /bin/passwd does not have the SUID bit (no 's' in the owner's execute field). As a result, user "test" receives an "Authentication token manipulation error". The password can be changed as root, which confirms it's a permissions/SUID issue.
Other options:
* B. If the password didn??t meet requirements, a different error would appear.
* C. There is no minimum day limit preventing password change (see chage -l output).
* D. The account and password are active (not disabled).
Reference:
CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 6: "User and Group Management", Section: "Managing User Passwords and Policies"
CompTIA Linux+ XK0-006 Objectives, Domain 1.0: System Management

NEW QUESTION 90
.....

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