CompTIA

Exam Questions XK0-006

CompTIA Linux+ Exam

NEW QUESTION 1
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 2
Which of the following commands should an administrator use to see a full hardware inventory of a Linux system?

A. dmidecode
B. lsmod
C. dmesg
D. lscpu

Answer: A

Explanation:
Hardware inventory and system information gathering are core responsibilities in Linux system management and are explicitly covered in CompTIA Linux+ V8 objectives. Among the listed commands, dmidecode is the most comprehensive tool for retrieving detailed hardware inventory information.
The dmidecode command reads data directly from the system??s DMI (Desktop Management Interface) / SMBIOS tables, which are provided by the system firmware (BIOS or UEFI). It reports detailed information about system hardware components, including motherboard details, BIOS version, system manufacturer, CPU sockets, memory slots, installed RAM modules, serial numbers, and asset tags. This makes it the preferred tool when a full hardware inventory is required.
The other options provide only partial or specific information. lsmod lists currently loaded kernel modules and does not provide physical hardware inventory. dmesg displays kernel ring buffer messages, which may include hardware detection logs but are not structured or complete inventory data. lscpu reports CPU architecture and processor details only, not the entire system hardware.
Linux+ V8 documentation highlights dmidecode as the authoritative utility for system hardware discovery and inventory auditing. It is commonly used in enterprise environments for documentation, troubleshooting, capacity planning, and compliance reporting.
Because it provides the most complete and authoritative hardware information available from the system firmware, the correct answer is A. dmidecode.

NEW QUESTION 3
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 4
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 5
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 6
While hardening a system, an administrator runs a port scan with Nmap, which returned the following output:

```
# nmap 104.21.75.76
Starting Nmap 7.80 ( https://nmap.org ) at 2024-07-09 18:09 UTC
Nmap scan report for 104.21.75.76
Host is up (0.00087s latency).
Not shown: 996 closed ports
PORT STATE SERVICE
23/tcp open telnet
80/tcp open http
443/tcp open https
8080/tcp open http-proxy

Nmap done: 1 IP address (1 host up) scanned in 4.93 seconds
```

Which of the following is the best way to address this security issue?

A. Configuring a firewall to block traffic on port 23 on the server
B. Changing the system administrator's password to prevent unauthorized access
C. Closing port 80 on the network switch to block traffic
D. Disabling and removing the Telnet service on the server

Answer: D

Explanation:
This scenario falls under theSecuritydomain of the CompTIA Linux+ V8 objectives and focuses on system hardening and service minimization. The Nmap scan output reveals thatport 23 (Telnet)is open on the system, which represents a significant security risk.
Telnet is an insecure, legacy protocol that transmits authentication credentials and session data inplaintext, making it vulnerable to interception through packet sniffing or man-in-the-middle attacks. Linux+ V8 documentation explicitly emphasizes the principle ofleast functionality, which states that unnecessary or insecure services should be disabled and removed entirely rather than merely restricted.
OptionD, disabling and removing the Telnet service on the server, is the best and most secure solution. This action eliminates the vulnerable service completely, ensuring that it cannot be exploited internally or externally. In secure Linux environments, Telnet should be replaced withSSH, which provides encrypted communication and strong authentication mechanisms.

OptionA, blocking port 23 with a firewall, reduces exposure but does not eliminate the underlying risk. If the firewall rules are misconfigured or bypassed, the Telnet service would still be available. Linux+ V8 best practices recommend removing insecure services rather than relying solely on perimeter controls.
OptionBis unrelated, as changing passwords does not address the risk of plaintext credential transmission. OptionCis incorrect because closing ports at the network switch level is not an appropriate or scalable solution for host-level service hardening and does not address internal access risks.
Linux+ V8 documentation consistently highlights service auditing, port scanning, and removal of insecure protocols as essential system hardening steps.
Therefore, the most effective and secure remediation is to disable and remove the Telnet service.

NEW QUESTION 7
A Linux systems administrator is running an important maintenance task that consumes a large amount of CPU, causing other applications to slow. Which of the following actions should the administrator take to help alleviate the issue?

A. Increase the available CPU time with pidstat.
B. Lower the priority of the maintenance task with renice.
C. Run the maintenance task with nohup.
D. Execute the other applications with the bg utility.

Answer: B

Explanation:
Process scheduling and resource management are essential Linux administration skills covered in Linux+ V8. When a process consumes excessive CPU resources, it can negatively impact overall system performance.
The correct solution is tolower the priorityof the CPU-intensive task using the renice command. Niceness values influence how much CPU time a process receives relative to others. Increasing the niceness value reduces the process??s priority, allowing other applications to receive CPU resources more fairly.
OptionBdirectly addresses the issue. The other options do not. pidstat monitors processes but does not modify CPU allocation. nohup allows a process to continue running after logout but does not affect scheduling priority. bg resumes a stopped job in the background but does not reduce CPU usage.
Linux+ V8 documentation explicitly references nice and renice for managing CPU contention. Therefore, the correct answer isB.

NEW QUESTION 8
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 9
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 10
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 10
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 11
An administrator is trying to terminate a process that is not responding. Which of the following commands should the administrator use in order to force the termination of the process?

A. kill PID
B. kill -1 PID
C. kill -9 PID
D. kill -15 PID

Answer: C

Explanation:
Comprehensive and Detailed Explanation From Exact Extract:
The kill command is used to send signals to processes. The -9 option sends the SIGKILL signal, which immediately terminates the process and cannot be caught or ignored by the process. This is used as a last resort when a process is not responding to the default (SIGTERM, -15) or other signals. The SIGKILL signal guarantees termination.
Other options:
* A.Default kill sends SIGTERM (-15), which requests a graceful shutdown but can be ignored.
* B.-1 sends SIGHUP, used to reload configuration, not terminate.
* D.-15 sends SIGTERM, not guaranteed to kill an unresponsive process.
[Reference:, CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 3: "Managing Processes", Section: "Sending Signals to Processes", CompTIA Linux+ XK0-006 Objectives, Domain 1.0: System Management, ,]

NEW QUESTION 14
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 /.

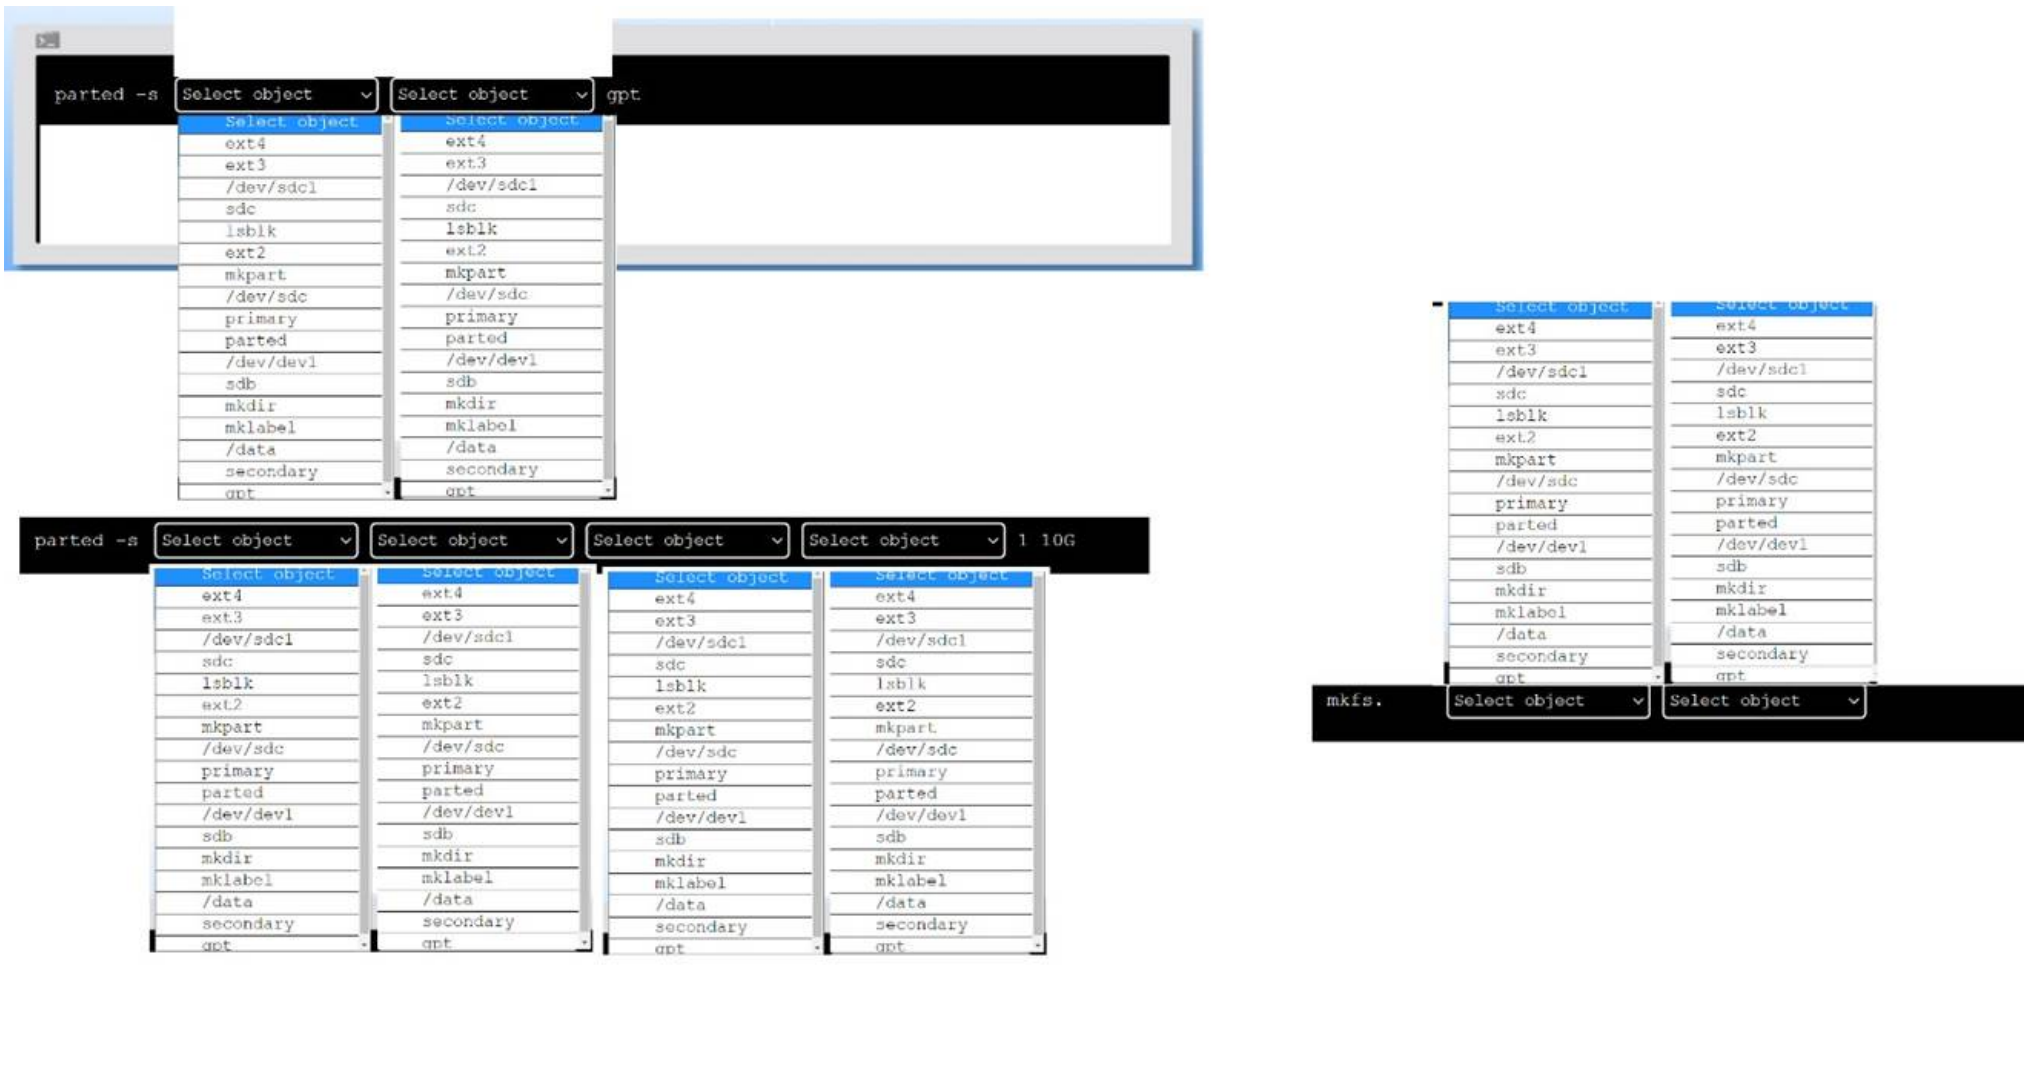

A. Mastered
B. Not Mastered

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 17
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 21
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 22
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 25
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 29
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 30
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 34
Which of the following filesystems contains non-persistent or volatile data?

A. /boot
B. /usr
C. /proc
D. /var

Answer: C

Explanation:
Understanding Linux filesystems and their purposes is a fundamental system management skill outlined in the Linux+ V8 objectives. Among the listed options,/procis the filesystem that contains non-persistent, volatile data.
The /proc filesystem is avirtual filesystemthat exists entirely in memory and is dynamically generated by the Linux kernel. It does not store data on disk and does not persist across system reboots. Instead, /proc provides real-time information about running processes, kernel parameters, system memory, CPU statistics, and hardware state. Files within /proc represent kernel data structures and change constantly as the system operates.
The other filesystems contain persistent data stored on disk. /boot stores bootloader files and kernel images, which are critical for system startup. /usr contains user applications, libraries, and documentation, all of which are persistent. /var holds variable data such as logs, spool files, and caches, which may change frequently but are still stored persistently on disk.
Linux+ V8 documentation emphasizes that /proc is used primarily for system monitoring and tuning. Administrators often interact with /proc to inspect process details or modify kernel parameters using tools like sysctl. Because its contents are generated at runtime and cleared on reboot, /proc is classified as non-persistent or volatile.
Therefore, the correct answer isC. /proc.

NEW QUESTION 36
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 37
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 39
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 43
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 47
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 51
An administrator logs in to a Linux server and notices the clock is 37 minutes fast. Which of the following commands will fix the issue?

A. hwclock
B. ntpdate
C. timedatectl
D. ntpd -q

Answer: B

Explanation:
Comprehensive and Detailed Explanation From Exact Extract:
The ntpdate command synchronizes the system clock with a remote NTP server immediately, correcting any significant time drift. This is ideal for one-time corrections.
For example:
bash
CopyEdit
ntpdate pool.ntp.org
Other options:
* A.hwclock reads or sets the hardware clock, but does not sync with network time.
* C.timedatectl can set the time manually or manage time settings, but does not immediately sync with a remote NTP server.
* D.ntpd -q can also sync the clock once, but ntpdate is designed specifically for immediate synchronization and is more straightforward for one-time corrections.
[Reference:, CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 5: "System Management", Section: "Time Synchronization", CompTIA Linux+ XK0-006 Objectives, Domain 1.0: System Management, ==========]

NEW QUESTION 52
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 53
A Linux administrator needs to securely erase the contents of a hard disk. Which of the following commands is the best for this task?

A. sudo rm -rf /dev/sda1
B. sudo shred /dev/sda1
C. sudo parted rm /dev/sda1
D. sudo dd if=/dev/null of=/dev/sda1

Answer: B

Explanation:
Secure data destruction is an important security requirement addressed in Linux+ V8 objectives. When data must be permanently erased, standard file deletion commands are insufficient because they do not overwrite the data on disk.
The shred command is specifically designed to securely erase files or block devices by overwriting them multiple times with random data. Using sudo shred /dev/sda1 overwrites the entire partition, making data recovery extremely difficult or impossible. This aligns directly with Linux+ V8 best practices for secure data sanitization.
The other options are incorrect. rm -rf removes directory entries but does not overwrite disk data. parted rm deletes partition entries but leaves the underlying data intact. dd if=/dev/null writes zero bytes and does not overwrite existing data blocks.
Linux+ V8 documentation identifies shred as the most appropriate tool for secure erasure when compliance or confidentiality is required. Therefore, the correct answer isB.

NEW QUESTION 57
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 58
A systems administrator is configuring new Linux systems and needs to enable passwordless authentication between two of the servers. Which of the following

commands should the administrator use?

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

Answer: A

NEW QUESTION 60
Which of the following best describes the role of initrd?

A. It is required to connect to the system via SSH.
B. It contains basic kernel modules and drivers required to start the system.
C. It contains trusted certificates and secret keys of the system.
D. It is required to initialize a random device within a Linux system.

Answer: B

NEW QUESTION 64
.....

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