Exam Questions XK0-006

CompTIA Linux+ Exam

https://www.2passeasy.com/dumps/XK0-006/

NEW QUESTION 1
A Linux administrator receives reports about MySQL service availability issues. The administrator observes the following information:

- uptime -p shows the system has been up for only 2 minutes
- journalctl shows messages indicating:mysqld invoked oom-killermysqld cpuset=/ mems_allowed=0 Which of the following explains why the server was offline?

A. The process exhausted server memory.
B. The process was intentionally terminated by a privileged user.
C. The process crashed because of a filesystem error.
D. A network outage caused a service availability issue.

Answer: A

Explanation:
is A. The process exhausted server memory.

NEW QUESTION 2
A systems administrator needs to integrate a new storage array into the company's existing storage pool. The administrator wants to ensure that the server is able to detect the new storage array. Which of the following commands should the administrator use to ensure that the new storage array is presented to the systems?

A. lsscsi
B. lsusb
C. lsipc
D. lshw

Answer: A

Explanation:
Comprehensive and Detailed Explanation: From Exact Extract:
The lsscsi command is used to list information about SCSI devices (including storage arrays) that are attached to the system. This is critical when integrating a new storage array because it allows the administrator to verify that the operating system detects the new device at the SCSI layer, which is the underlying interface for most enterprise storage solutions. lsscsi outputs a list of recognized SCSI devices, their device nodes, and associated information.
Other options:

- B. lsusb: Lists USB devices, not storage arrays on SCSI/SATA/SAS.
- C. lsipc: Displays information on IPC (inter-process communication) facilities, unrelated to hardware detection.
- D. lshw: Lists hardware details and can show storage, but lsscsi is specifically designed for SCSI device detection and is the most direct method for this task.

Reference:
CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 7: "Managing Storage", Section: "Identifying and Accessing Storage Devices"
CompTIA Linux+ XK0-006 Objectives: Domain 4.0 – Storage and Filesystems
==========

NEW QUESTION 3
A systems administrator is reconfiguring existing user accounts in a Linux system. Which of the following commands should the administrator use to include "myuser" in the finance group?

A. groupadd finance myuser
B. groupmod finance myuser
C. useradd -g finance myuser
D. usermod -aG finance myuser

Answer: D

Explanation:
Comprehensive and Detailed Explanation: From Exact Extract:
To add an existing user (myuser) to an existing group (finance) without removing them from other groups, the correct command is usermod -aG finance myuser.
The -aG option appends the user to the supplementary group
(s) specified.
Other options:

- A. groupadd is for creating new groups, not adding users to groups.
- B. groupmod is for modifying group properties, not user membership.
- C. useradd creates new users; not applicable to existing users.

Reference:
CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 6: "User and Group Management", Section: "Modifying Group Membership"
CompTIA Linux+ XK0-006 Objectives, Domain 1.0: System Management
==========

NEW QUESTION 4
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 5
A systems administrator needs to enable routing of IP packets between network interfaces. Which of the following kernel parameters should the administrator change?

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

Answer: D

Explanation:
IP packet forwarding is a key networking function in Linux system management and is explicitly referenced in the Linux+ V8 objectives. Enabling this feature allows a Linux system to act as a router by forwarding packets between network interfaces.
The kernel parameter responsible for this behavior is net.ipv4.ip_forward. When this parameter is set to 1, the Linux kernel allows IPv4 packets to be forwarded between interfaces. By default, this setting is often disabled on non-routing systems for security reasons.
The parameter can be modified temporarily using the sysctl command or permanently by editing /etc/sysctl. conf or files under /etc/sysctl.d/. Linux+ V8 documentation highlights this parameter as essential for configuring routing, NAT, and firewall-based gateway systems.
The other options are incorrect. net.ipv4.ip_multicast controls multicast behavior, not packet forwarding. net. ipv4.ip_route is not a valid kernel parameter. net.ipv4.ip_local_port_range defines the range of ephemeral ports used by outgoing connections and has no effect on routing.
Properly enabling IP forwarding is critical when configuring VPN gateways, firewalls, and network appliances. Therefore, the correct answer is D. net.ipv4.ip_forward.

NEW QUESTION 6
A systems administrator is creating a backup copy of the /home/ directory. Which of the following commands allows the administrator to archive and compress the directory at the same time?

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

Answer: C

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

NEW QUESTION 7
Which of the following 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 8
A systems administrator manages multiple Linux servers and needs to set up a reliable and secure way to handle the complexity of managing event records on the OS and application levels. Which of the following should the administrator do?

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

Answer: B

Explanation:
Log management is a critical system management function highlighted in CompTIA Linux+ V8, particularly in multi-server environments. As the number of systems and applications grows, managing logs locally on each server becomes inefficient and error-prone.
The best solution is toimplement a centralized log aggregation solution, making optionBcorrect. Centralized logging collects logs from multiple systems and applications into a single, secure location. This simplifies monitoring, searching, correlation, auditing, and incident response. Common solutions include syslog servers, ELK/EFK stacks, and SIEM platforms.
Linux+ V8 documentation emphasizes centralized logging as a best practice for availability, troubleshooting, and security analysis. It enables administrators to detect patterns, investigate incidents, and maintain compliance more effectively than isolated log files.
The other options are insufficient on their own. On-demand retrieval does not scale well. Log backups protect data but do not simplify analysis. Log rotation manages disk usage but does not address distributed log complexity.
Therefore, the correct answer isB. Implement a centralized log aggregation solution.

NEW QUESTION 9
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 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 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 14
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 16
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 18
A Linux user needs to download the latest Debian image from a Docker repository. Which of the following commands makes this task possible?

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

Answer: B

Explanation:
Container management and image handling are part of modern Linux automation practices covered in CompTIA Linux+ V8. Docker images are stored in container registries such as Docker Hub, and administrators commonly need to download images to deploy containers.
The correct command for downloading an image from a Docker repository is docker image pull. This command retrieves the specified image from a configured container registry and stores it locally. When no tag is specified, Docker automatically pulls the latest available version of the image. Therefore, docker image pull debian downloads the most recent Debian image from Docker Hub.
The other options are incorrect. docker image init is not a valid Docker command and does not exist in Docker??s CLI. docker image import is used to create a Docker image from a tarball file, not to download an image from a repository. docker image save exports an existing local image into a tar archive and does not retrieve images from a remote registry.
Linux+ V8 documentation emphasizes understanding container image lifecycles, including pulling, tagging, and running images. Pulling images is a foundational step before container execution and automation workflows.
Therefore, the correct answer is B. docker image pull debian.

NEW QUESTION 19
A DevOps engineer made some changes to files in a local repository. The engineer realizes that the changes broke the application and the changes need to be reverted back. Which of the following commands is the best way to accomplish this task?

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

Answer: B

Explanation:
Version control rollback operations are a core DevOps skill covered in the Linux+ V8 objectives. When changes in a local Git repository break an application and must be reverted, the administrator must choose a command that directly undoes those changes.
The command git reset is the most appropriate option in this scenario. It allows the engineer to move the current branch pointer (HEAD) to a previous commit, effectively discarding or undoing local changes. Depending on the reset mode (--soft, --mixed, or --hard), the engineer can control whether changes are preserved in the staging area or working directory. This flexibility makes git reset the primary tool for reverting problematic local changes.
The other options are not suitable. git pull fetches and merges changes from a remote repository and does not revert local modifications. git rebase rewrites commit history and is used to reapply commits on top of another base, not to undo broken changes. git stash temporarily saves uncommitted changes for later use but does not revert the repository to a stable state.
Linux+ V8 documentation emphasizes that git reset is commonly used during local development when changes need to be undone quickly before being shared with others. Therefore, the correct answer is B.

NEW QUESTION 20
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 23
Users cannot access a server after it has been restarted. At the server console, the administrator runs the following commands;

```
$ ss -lnt
State      Recv-Q Send-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 27
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 32
Which of the following passwords is the most complex?

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

Answer: C

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

NEW QUESTION 33
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 38
A Linux administrator wants to add a user to the Docker group without changing the user??s primary group. Which of the following commands should the administrator use to complete this task?

A. sudo groupmod docker user
B. sudo usermod -g docker user
C. sudo usermod -aG docker user
D. sudo groupmod -G docker user

Answer: C

Explanation:
User and group management is a core System Management topic in CompTIA Linux+ V8. When adding a user to an additional group—such as the docker group—care must be taken not to alter the user's primary group.
The correct command issudo usermod -aG docker user. The -G option specifies a supplementary group, and the -a (append) option ensures the user is added to the groupwithout removing existing group memberships. This is especially important because omitting -a would overwrite the user's supplementary groups.
OptionB, usermod -g docker user, changes the user'sprimary group, which is not desired. OptionsAandDmisuse groupmod, which is intended for modifying group properties, not user membership.
Linux+ V8 documentation explicitly warns that failing to use -a with -G can unintentionally remove a user from all other supplementary groups, potentially causing access issues.
Therefore, the correct and safe command isC. sudo usermod -aG docker user.

NEW QUESTION 40
A Linux administrator needs to analyze a compromised disk for traces of malware. To complete the analysis, the administrator wants to make an exact, block-level copy of the disk. Which of the following commands accomplishes this task?

A. cp -rp /dev/sdc/* /tmp/image
B. cpio -i /dev/sdc -ov /tmp/image
C. tar cvzf /tmp/image /dev/sdc
D. dd if=/dev/sdc of=/tmp/image bs=8192

Answer: D

Explanation:
Disk forensics and malware analysis fall under theSecuritydomain in the CompTIA Linux+ V8 objectives. When analyzing a compromised disk, it is critical to preserve the data exactly as it exists, including unused space, deleted files, and hidden metadata. This requires ablock-level copy, not a file-level copy.
The dd command is the correct tool for this task. It operates at a low level, copying raw data from an input device (if=/dev/sdc) directly to an output file (of=/tmp/image) without interpreting filesystem structures. This ensures an exact, bit-for-bit replica of the disk, which is essential for forensic integrity and malware analysis. The bs=8192 option improves performance by specifying a larger block size during copying.
The other options are incorrect. cp -rp copies files and directories but does not capture free space, deleted data, or disk metadata. cpio and tar are archive utilities that operate at the filesystem level and cannot produce a true disk image. These tools also require the filesystem to be mounted and readable, which is not appropriate for forensic preservation.
Linux+ V8 documentation highlights dd as the preferred utility for disk imaging, backups, and forensic investigations. Administrators are also advised to perform such operations on unmounted disks to avoid altering evidence.
Therefore, the correct and best command for creating an exact block-level disk copy isD. dd if=/dev/sdc of=/tmp/image bs=8192.

NEW QUESTION 41
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 46
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 51
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 53
An administrator receives reports that a web service is not responding. The administrator reviews the following outputs:

```
$ echo $PWD
/etc/pki/nginx

$ ls -lRt
.:
total 8
drwxr-xr-x. 2 root root 6 Jul 10 10:57 private
-rw-r--r--. 1 root root 895 Jul 10 10:56 server.crt
-rw-------. 1 root root 227 Jul 10 10:56 server.key
./private:
total 0

$ sudo systemctl status nginx
  nginx.service - The nginx HTTP and reverse proxy server
       Loaded: loaded (/usr/lib/systemd/system/nginx.service; disabled; preset: disabled)
       Active: failed (Result: exit-code) since Wed 2023-11-01 06:56:51 EDT; 6s ago
    Process: 110551 ExecStartPre=/usr/bin/rm -f /run/nginx.pid (code=exited, status=0/SUCCESS)
    Process: 110552 ExecStartPre=/usr/sbin/nginx -t (code=exited, status=1/FAILURE)
        CPU: 144ms

Nov 01 06:56:51 webserver systemd[1]: Starting nginx.service - The nginx HTTP and reverse proxy server...
Nov 01 06:56:51 webserver nginx[110552]: nginx: [emerg] cannot load certificate key "/etc/pki/nginx/private/server.key": BIO_new_file()
failed (SSL: error:80000002:system library::No such file or directory:calli>
Nov 01 06:56:51 webserver nginx[110552]: nginx: configuration file /etc/nginx/nginx.conf test failed
Nov 01 06:56:51 webserver systemd[1]: nginx.service: Control process exited, code=exited, status=1/FAILURE
Nov 01 06:56:51 webserver systemd[1]: nginx.service: Failed with result 'exit-code'.
Nov 01 06:56:51 webserver systemd[1]: Failed to start nginx.service - The nginx HTTP and reverse proxy server.
```

Which of the following is the reason the web service is not responding?

A. The private key needs to be renamed from server.crt to server, key so the service can find it.
B. The private key does not match the public key, and both keys should be replaced.
C. The private key is not in the correct location and needs to be moved to the correct directory.
D. The private key has the incorrect permissions and should be changed to 0755 for the service.

Answer: C

Explanation:
This issue falls under the Troubleshooting domain of the CompTIA Linux+ V8 objectives, specifically service startup failures and certificate-related errors. The provided output clearly indicates that the NGINX service fails during startup due to an inability to locate the private key file.
The critical error message is:
cannot load certificate key "/etc/pki/nginx/private/server.key": No such file or directory
This message confirms that NGINX is explicitly configured to look for the private key in the directory /etc/pki/nginx/private/. However, the directory listing shows that the private directory exists but is empty, while the server.key file is located in /etc/pki/nginx/ instead. Because NGINX cannot find the private key at the configured path, the configuration test (nginx -t) fails, and systemd prevents the service from starting.
Option C correctly identifies the root cause: the private key is not in the correct location. Moving server.key into /etc/pki/nginx/private/ (or updating the NGINX configuration to match the current location) would resolve the issue. Linux+ V8 documentation stresses that service failures often result from misaligned configuration paths rather than corrupted files.
The other options are incorrect. Option A incorrectly refers to renaming a certificate file and does not address the path issue. Option B suggests a key mismatch, which would generate a different SSL error rather than a "file not found" error. Option D is also incorrect because private keys should not have executable permissions like 0755; typically, they are restricted (for example, 0600) for security reasons.
Therefore, the web service is not responding because the private key file is not located in the directory expected by the NGINX configuration. The correct answer is C.

NEW QUESTION 57
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 58
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 59
An administrator wants to search a file named myFile and look for all occurrences of strings containingat least five characters, wherecharacters two and five are i, butcharacter three is not b. Which of the following commands should the administrator execute to get the intended result?

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

Answer: D

Explanation:
Pattern matching using regular expressions is a key troubleshooting and text-processing skill covered in CompTIA Linux+ V8. The grep command, combined with regular expressions, allows administrators to search for complex string patterns within files.
The requirement specifies:
The string must containat least five characters
Character 2must be i
Character 3mustnotbe b
Character 5must be i
To meet these conditions, the correct regular expression structure is:
. ?? any character (position 1)
i ?? literal i (position 2)
[^b] ?? any character except b (position 3)
. ?? any character (position 4)
i ?? literal i (position 5)
This results in the expression:
i[^b].i
OptionD, grep .i[^b].i myFile, correctly implements this logic. It ensures positional matching and excludes unwanted characters using a negated character class ([^b]), which is explicitly covered in Linux+ V8 regular expression objectives.
The other options contain invalid or malformed regular expressions and do not meet the positional or exclusion requirements. Linux+ V8 emphasizes understanding anchors, character classes, and position-based matching when troubleshooting log files or configuration data.
Therefore, the correct answer isD.

NEW QUESTION 64
Which of the following is a reason multiple password changes on the same day are not allowed?

A. To avoid brute-forced password attacks by making them too long to perform
B. To increase password complexity and the system's security
C. To stop users from circulating through the password history to return to the originally used password
D. To enforce using multifactor authentication with stronger encryption algorithms instead of passwords

Answer: C

Explanation:
Password policy enforcement is a critical component of system security covered in the CompTIA Linux+ V8 objectives. One common control implemented in Linux systems is restricting how frequently users can change their passwords, often referred to asminimum password ageenforcement.
The primary reason multiple password changes within a short time frame are not allowed is toprevent password cycling attacks. Without this restriction, a user could repeatedly change their password in quick succession to bypass password history controls and eventually reuse a previously compromised or weak

password. OptionCaccurately describes this scenario and aligns directly with Linux+ V8 security guidance.
Linux systems enforce this behavior through tools such as chage and PAM (Pluggable Authentication Modules). Administrators can configure minimum password age values to ensure users must wait a defined period before changing passwords again. This ensures that password history requirements are effective and meaningful.
The other options are incorrect. OptionAconfuses password expiration with brute-force mitigation, which is typically addressed through account lockout policies. OptionBrefers to password complexity, which is enforced through character requirements rather than change frequency. OptionDis unrelated, as password expiration policies do not enforce multifactor authentication.
Linux+ V8 documentation emphasizes layered access controls, and preventing password reuse through enforced timing restrictions is a core principle of secure authentication design.
Therefore, the correct answer isC.

NEW QUESTION 65
Which of the following cryptographic functions ensures a hard drive is encrypted when not in use?

A. GPG
B. LUKS
C. PKI certificates
D. OpenSSL

Answer: B

Explanation:
Disk encryption is a key Linux+ V8 security objective, especially for protecting data at rest.LUKS (Linux Unified Key Setup)is the standard Linux framework for full-disk and partition-level encryption.
OptionBis correct. LUKS provides strong encryption for storage devices, ensuring that data remains unreadable when the system is powered off or the disk is removed. It integrates with tools like cryptsetup and supports key management, passphrases, and multiple unlock methods.
The other options are incorrect. GPG encrypts files, not entire disks. PKI certificates are used for identity and trust, not disk encryption. OpenSSL is a cryptographic library, not a disk encryption mechanism.
Linux+ V8 documentation explicitly identifies LUKS as the primary solution for disk encryption on Linux systems. Therefore, the correct answer isB.

NEW QUESTION 68
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual XK0-006 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the XK0-006 Product From:

https://www.2passeasy.com/dumps/XK0-006/

Money Back Guarantee

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