CompTIA

Exam Questions XK0-006

CompTIA Linux+ Exam

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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
On a Kubernetes cluster, which of the following resources should be created in order to expose a port so it is publicly accessible on the internet?

A. Deployment
B. Network
C. Service
D. Pod

Answer: C

Explanation:
Container orchestration concepts are part of the Automation and Orchestration domain in Linux+ V8. In Kubernetes, workloads run inside Pods, but Pods are not directly accessible from outside the cluster.
To expose an application externally, aServiceresource must be created. Services provide a stable network endpoint and can be configured as NodePort, LoadBalancer, or ClusterIP. Public exposure is typically achieved using NodePort or LoadBalancer types.
OptionC, Service, is correct. Deployments manage Pods, but they do not handle networking exposure. Pods represent running containers but lack external accessibility by default. "Network" is not a valid Kubernetes resource type.
Linux+ V8 documentation highlights Services as the mechanism for exposing containerized applications. Therefore, the correct answer isC.

NEW QUESTION 3
Which of the following is the main reason for setting up password expiry policies?

A. To avoid using the same passwords repeatedly
B. To mitigate the use of exposed passwords
C. To force usage of passwordless authentication
D. To increase password strength and complexity

Answer: B

Explanation:
Password management is a core topic in the Security domain of CompTIA Linux+ V8. Password expiry policies are implemented to reduce the risk associated with long-lived credentials.
The primary reason for enforcing password expiration is to mitigate the risk of exposed or compromised passwords. If a password is leaked through phishing, malware, keylogging, or data breaches, limiting its lifespan reduces the window of opportunity for attackers to exploit it. Requiring periodic password changes ensures that compromised credentials eventually become invalid.
Option B correctly captures this security objective. Linux+ V8 documentation emphasizes minimizing credential exposure as a key principle of access control.
The other options are secondary or incorrect. Avoiding password reuse and increasing complexity are addressed through password history and complexity policies, not expiration alone. Password expiry does not force passwordless authentication, making option C incorrect.
Therefore, the correct answer is B. To mitigate the use of exposed passwords.

NEW QUESTION 4
A systems administrator attempts to edit a file as root, but receives the following error:

```
E212: Cannot open file for writing

# ls -l /etc/resolv.conf
-rw-------. 1 root admin 141 May 30 11:00 /etc/resolv.conf

# lsattr /etc/resolv.conf
----i----------- /etc/resolv.conf
```

Which of the following commands allows the administrator to edit the file?

A. chown root /etc/resolv.conf
B. chattr -i /etc/resolv.conf
C. chmod 750 /etc/resolv.conf
D. chgrp root /etc/resolv.conf

Answer: B

Explanation:
This scenario involves Linux file attributes and falls under theSystem Managementdomain of the CompTIA Linux+ V8 objectives. Although the administrator is operating as the root user, the system prevents the file from being modified. This behavior indicates that standard UNIX permissions are not the root cause of the problem.
The critical clue is provided by the lsattr /etc/resolv.conf output, which shows theimmutable (i) attributeset on the file. When a file is marked immutable, it cannot be modified, deleted, renamed, or written to by any user, including root. This restriction overrides normal file permissions and ownership settings.
The chattr command is used to modify extended file attributes on Linux filesystems such as ext4. The option -i specifically removes the immutable attribute, restoring the file??s ability to be edited. Therefore, running chattr -i /etc/resolv.conf allows the administrator to open and modify the file successfully.
The other options do not resolve the issue. chown root changes file ownership, but the file is already owned by root. chmod 750 modifies permission bits, but permissions are ignored when the immutable attribute is set. chgrp root changes the group ownership, which also has no effect when immutability is enforced.
Linux+ V8 documentation highlights immutable files as a security and stability feature, commonly used to protect critical configuration files from accidental or unauthorized changes. Administrators must explicitly remove this attribute before making modifications.
Therefore, the correct command to allow editing the file isB. chattr -i /etc/resolv.conf.

NEW QUESTION 5
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 6
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 7
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 8
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 9
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 /.

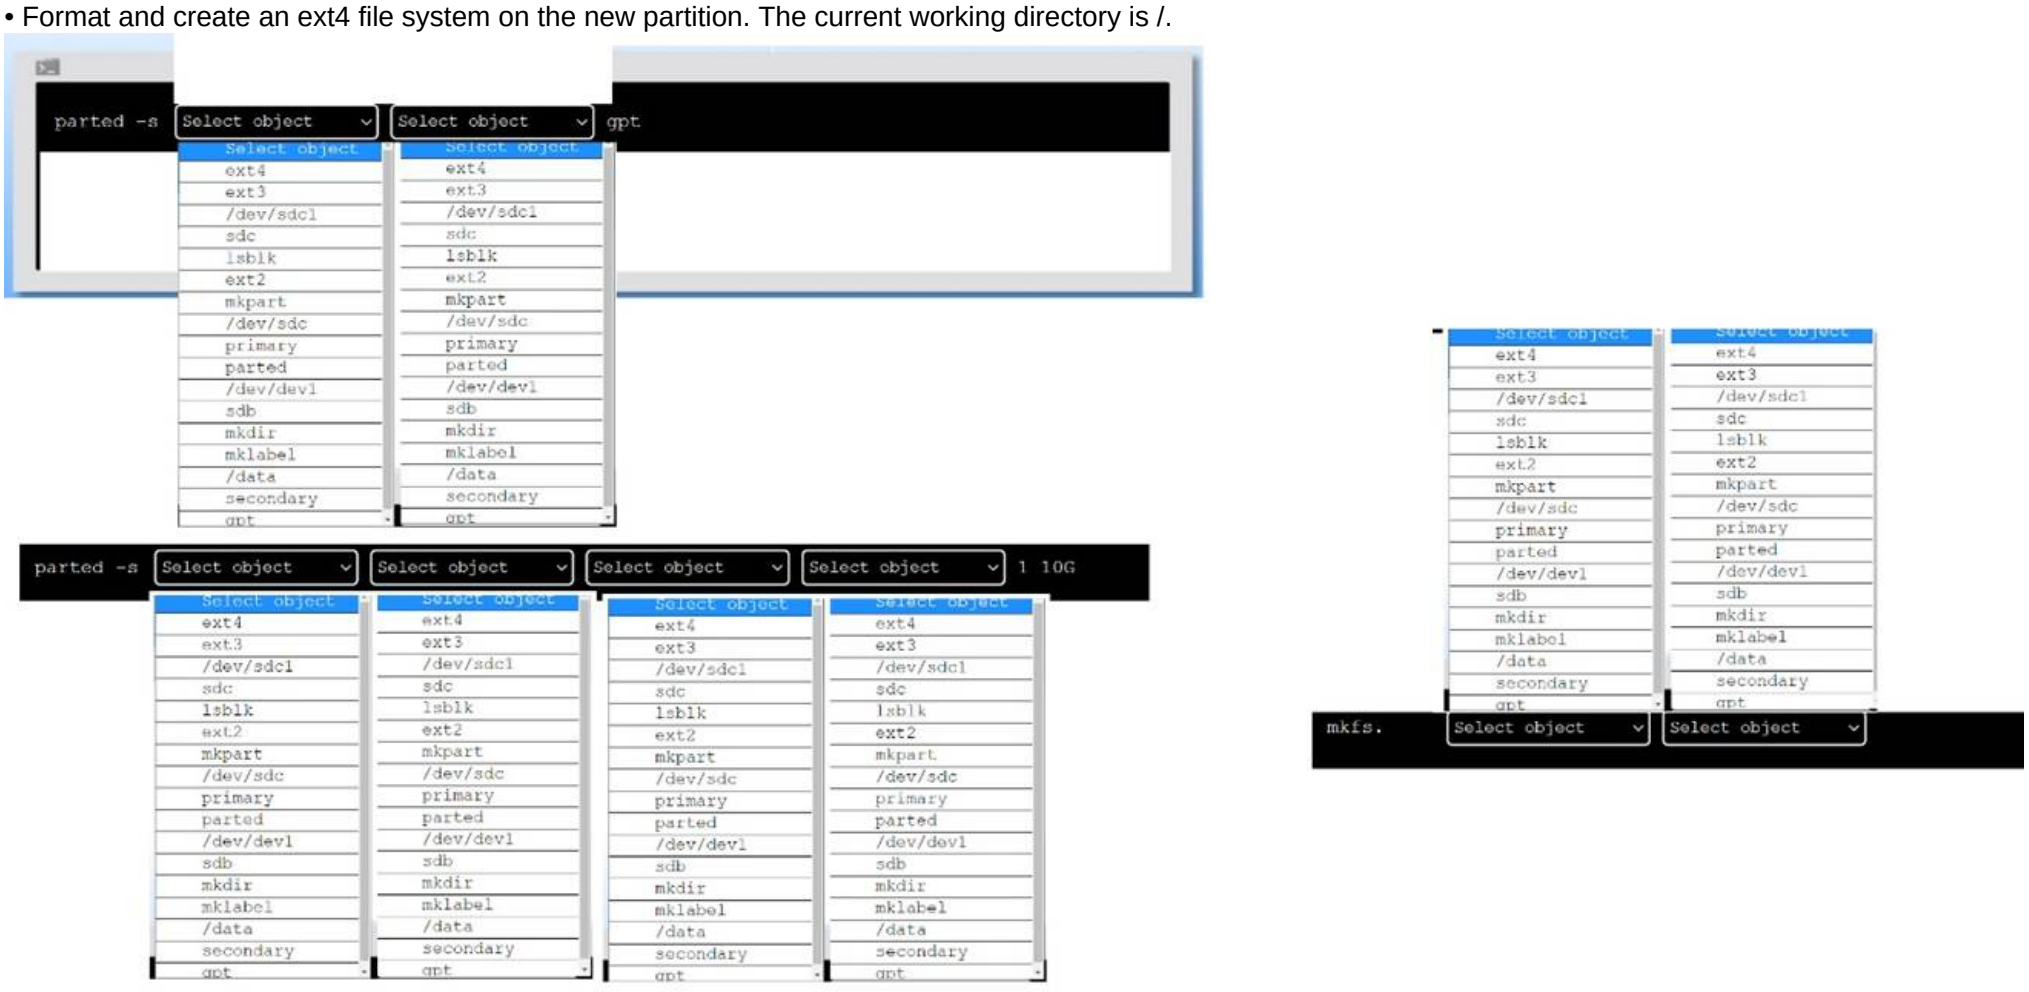

A. Mastered
B. Not Mastered

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 10
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 10
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 12
Users cannot access a server after it has been restarted. At the server console, the administrator runs the following commands;

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

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

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

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

Which of the following is the cause of the issue?

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

Answer: B

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

NEW QUESTION 14
A systems administrator wants to review the logs from an Apache 2 error.log file in real time and save the information to another file for later review. Which of the following commands should the administrator use?

A. tail -f /var/log/apache2/error.log > logfile.txt
B. tail -f /var/log/apache2/error.log | logfile.txt
C. tail -f /var/log/apache2/error.log >> logfile.txt
D. tail -f /var/log/apache2/error.log | tee logfile.txt

Answer: D

Explanation:

Log monitoring is a common troubleshooting task in Linux system administration, and Linux+ V8 covers command-line tools for real-time log analysis. The requirement in this scenario is twofold: view log entries as they occur and simultaneously save them to another file.
The command tail -f /var/log/apache2/error.log | tee logfile.txt fulfills both requirements. The tail -f command follows the log file in real time, displaying new entries as they are written. The pipe (|) sends this output to the tee command, which writes the data to logfile.txt while also displaying it on standard output.
The other options are insufficient. Option A redirects output to a file but prevents real-time viewing. Option C appends output but still suppresses terminal display. Option B is syntactically invalid and does not use a proper command for writing output.
Linux+ V8 documentation specifically references tee as a useful utility for duplicating command output streams. This makes option D the correct and most effective solution.

NEW QUESTION 16
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 19
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 24
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 25
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 27
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 28
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 30
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 35
Which of the following best describes journald?

A. A system service that collects and stores logging data
B. A feature that creates crash dumps in case of kernel failure
C. A service responsible for keeping the filesystem journal

D. A service responsible for writing audit records to a disk

Answer: A

NEW QUESTION 39

.....

Relate Links

100% Pass Your XK0-006 Exam with Exambible Prep Materials

https://www.exambible.com/XK0-006-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/