CompTIA

Exam Questions XK0-006

CompTIA Linux+ Exam

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

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

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

Answer: A

Explanation:
is A. The process exhausted server memory.

NEW QUESTION 2
Which of the following is a protocol for accessing distributed directory services containing a hierarchy of users, groups, machines, and organizational units?

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

Answer: C

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

NEW QUESTION 3
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 4
In the echo "profile-$num-$name" line of a shell script, the variable $num seems to not be expanding during execution. Which of the following notations ensures the value is expanded?

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

Answer: D

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

Linux+ V8 documentation emphasizes ${VAR} notation as a best practice in shell scripting for clarity and correctness. Therefore, the correct answer is D.

NEW QUESTION 5
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 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
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 8
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 9
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 10
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 12
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 16
A junior system administrator removed an LVM volume by mistake.
INSTRUCTIONS
Part 1
Review the output and select the appropriate command to begin the recovery process.
Part 2
Review the output and select the appropriate command to continue the recovery process.
Part 3
Review the output and select the appropriate command to complete the recovery process and access the underlying data.

Part 1 | Part 2 | Part 3

>_ Commands

```
[root@comptiasim ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        1.9G     0  1.9G   0% /dev
tmpfs           1.9G     0  1.9G   0% /dev/shm
tmpfs           1.9G   17M  1.9G   1% /run
tmpfs           1.9G     0  1.9G   0% /sys/fs/cgroup
/dev/xvda1      8.0G  1.1G  7.0G  13% /
tmpfs           379M     0  379M   0% /run/user/1000
```

```
[root@comptiasim ~]# df -t
[root@comptiasim ~]# ls -l /dev | grep -v tty
[root@comptiasim ~]# ls -l /etc/lvm/archive
[root@comptiasim ~]# pvdisplay
[root@comptiasim ~]# pvs
[root@comptiasim ~]# vgcfgrestore --list vg01
[root@comptiasim ~]# vgdisplay
[root@comptiasim ~]# vgs
```

Select the appropriate command to begin the recovery process.

```
[root@comptiasim ~]#
```

Select command

- lvchange -a y /dev/vg01/lv01
- lvconvert --type mirror lv01
- pvscan
- vgcfgrestore vg01 -f /etc/lvm/archive/vg01_00002-966141411.vg
- vgcfgrestore vg01 -f /etc/lvm/backup/vg01
- lvchange -a n /dev/vg01/lv01
- vgcfgrestore vg01 -t -M /etc/lvm/archive/vg01_00001-810050352.vg

Select command

Part 2 | Part 2 | Part 2

>_ Commands

```
[root@comptiasim ~]# blkid
/dev/xvda1: UUID="388a99ed-9486-4a46-aeb6-06eaf6c47675" TYPE="xfs"
/dev/xvdf: UUID="1uyvyk-FfdO-8rvF-cYba-15zc-EHRZJN3Uhm" TYPE="LVM2 member"
```

```
[root@comptiasim ~]# blkid
[root@comptiasim ~]# dmesg | tail -20
[root@comptiasim ~]# blbik
[root@comptiasim ~]# ls /
[root@comptiasim ~]# lvdisplay
[root@comptiasim ~]# lvs
[root@comptiasim ~]# lvecan
[root@comptiasim ~]# pvs
[root@comptiasim ~]# vgs
```

Select the appropriate command to continue the recovery process.

```
[root@comptiasim ~]#
```

Select command

- pvchange -x y /dev/xvdf
- lvextend -L v54 vg01/lv01 /dev/xvdf
- lvchange -x y /dev/vg01/lvd1
- mount /dev/vg01/lv01/ /important data
- lvchange -a n /dev/vg01/lvd1

Select command

Part 1 | Part 2 | Part 3

Commands

```
[root@comptiasim ~]# blkid
[root@comptiasim ~]# cat /etc/fstap
[root@comptiasim ~]# ls -l /dev/napper/
[root@comptiasim ~]# ls -l /
[root@comptiasim ~]# lsblk
[root@comptiasim ~]# lvdisplay
[root@comptiasim ~]# lvscan
[root@comptiasim ~]# tail -f /var/log/messages
[root@comptiasim ~]# xfs_repair -n /dev/vg01/lv01
```

```
[root@comptiasim ~]# blkid
/dev/xvda1:              UUID="388a99ed-9486-4a46-aeb6-06eaf6c47675"
TYPE="xfs"
/dev/mapper/vg01-lv01:   UUID="c63883e9-ceca-45f4-9ad9-f8d8c1814e7e"
TYPE="xfs"
/dev/xvdf:               UUID="1uyyyk-Ffd0-RrYF-cYba-l5zC-EHRZ UN3UHm"
TYPE="LVM2_member"
```

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

```
[root@comptiasim ~]#  Select command
```

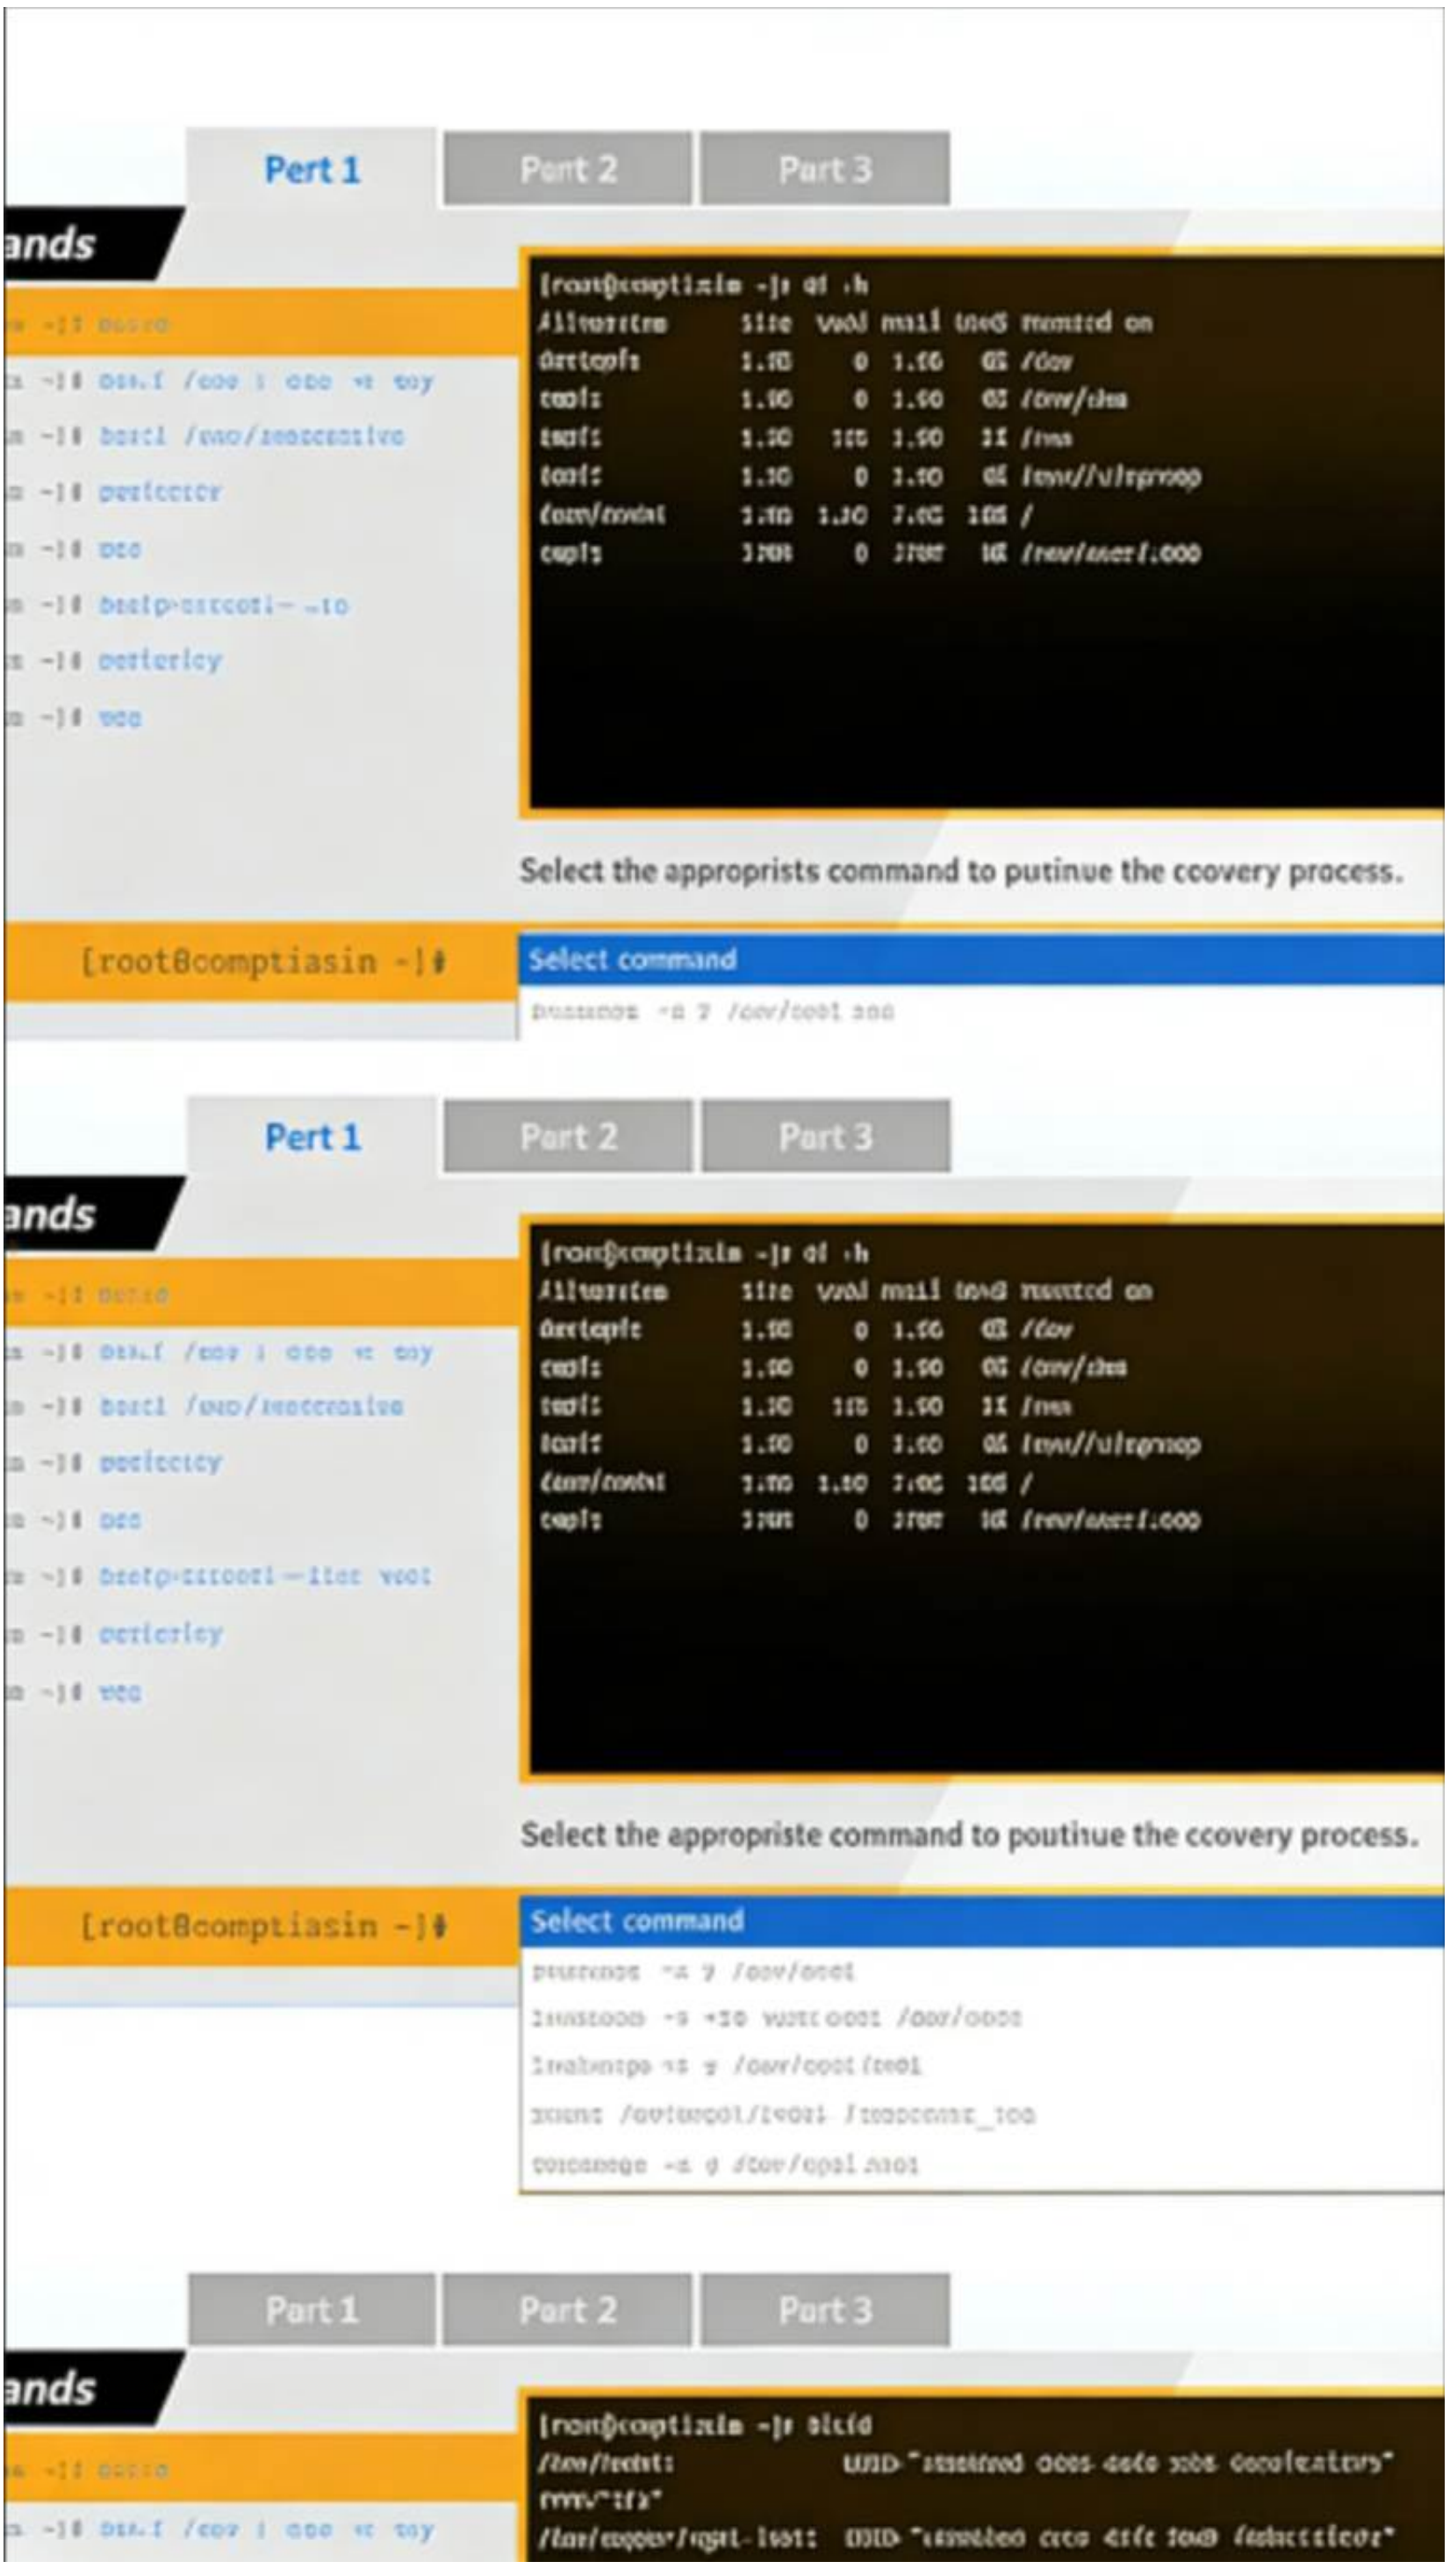
Part 1
Part 2
Part 3
Select the appropriste command to putinue the ccovery process.
[root@comptiasin ~]#
Select command
Part 1
Part 2
Part 3
Select the appropriste command to poutinue the ccovery process.
[root@comptiasin ~]#
Select command
Part 1
Part 2
Part 3

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Part 1 – Begin the recovery process Answer
vgcfgrestore vg01 -f /etc/lvm/archive/vg01_00001-810050352.vg
Part 2 – Continue the recovery process Answer
lvchange -ay /dev/vg01/lv01
Part 3 – Complete recovery and access data Answer
mount /dev/vg01/lv01 /important_data
This performance-based question tests LVM recovery, a critical System Management skill in CompTIA Linux+ V8. The scenario indicates that a logical volume was removed, but the underlying physical volume and volume group metadata still exist.
Part 1: Restoring Volume Group Metadata
The first screenshot shows that:
* Physical volumes (pvdisplay, pvs) still exist
* The logical volume is missing
* /etc/lvm/archive/ contains archived VG metadata
Linux automatically stores backups of LVM metadata in /etc/lvm/archive whenever changes are made. The correct first step is to restore the volume group metadata using:
vgcfgrestore vg01 -f /etc/lvm/archive/vg01_00001-810050352.vg
This restores the logical volume definitions but does not activate them yet.
This is the only correct starting point in Linux+ V8 recovery workflows.
Part 2: Activating the Logical Volume
After metadata restoration:
* The LV exists but is inactive
* blkid shows the LV as TYPE="LVM2_member"
The logical volume must be activated before it can be mounted: lvchange -ay /dev/vg01/lv01
This makes the LV available under /dev/vg01/lv01. Linux+ explicitly requires LV activation after recovery.
Part 3: Accessing the Data
The final output shows:
* The filesystem type is xfs
* The logical volume is now visible
Since there is no indication of filesystem corruption, no repair is required.
The correct final step is to mount the filesystem: mount /dev/vg01/lv01 /important_data
This restores full access to the underlying data.

NEW QUESTION 19
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 23
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 28
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 33
A Linux administrator wants to make the enable_auth variable set to 1 and available to the environment of subsequently executed commands. Which of the following should the administrator use for this task?

A. let ENABLE_AUTH=1
B. ENABLE_AUTH=1
C. ENABLE_AUTH=$(echo $ENABLE_AUTH)
D. export ENABLE_AUTH=1

Answer: D

Explanation:
Environment variables in Linux can exist either locally within a shell or be exported to child processes. CompTIA Linux+ V8 emphasizes the distinction betweenshell variablesandenvironment variables, as this affects how applications inherit configuration values.
OptionD, export ENABLE_AUTH=1, is the correct choice because it both assigns the variable and marks it for export to the environment. Once exported, the variable becomes available to all subsequently executed commands and child processes spawned from the current shell. This behavior is required when applications or scripts rely on environment variables for configuration.
OptionB, ENABLE_AUTH=1, only sets a shell-local variable. While it is accessible within the current shell session, it is not inherited by child processes unless explicitly exported. OptionA, let ENABLE_AUTH=1, performs arithmetic evaluation and does not export the variable. OptionCincorrectly assigns the output of a command substitution and does not set the desired value.
Linux+ V8 documentation highlights export as the correct mechanism for making variables available system-wide within a user session. Therefore, the correct answer isD.

NEW QUESTION 35
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 39
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 43
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 47
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 51
.....

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