CompTIA

Exam Questions XK0-006

CompTIA Linux+ Exam

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

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

Answer: A

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

NEW QUESTION 4
A systems administrator is preparing a Linux system for application setup. The administrator needs to create an environment variable with a persistent value in one of the user accounts. Which of the following commands should the administrator use for this task?

A. export "VAR=SomeValue" >> ~/.ssh/profile
B. export VAR=value
C. VAR=value
D. echo "export VAR=value" >> ~/.bashrc

Answer: D

Explanation:
Environment variables are widely used in Linux systems to configure application behavior, and Linux+ V8 emphasizes the distinction between temporary and persistent variables. A variable is persistent only if it is defined in a shell initialization file.
The correct approach is echo "export VAR=value" >> ~/.bashrc. This command appends the variable definition to the user??s .bashrc file, ensuring the variable is set automatically every time the user starts a new shell session. This makes the variable persistent for that specific user.
Options B and C only define variables in the current shell session and are lost when the session ends. Option A incorrectly targets the SSH configuration directory and is not appropriate for defining shell environment variables.
Linux+ V8 documentation highlights .bashrc, .bash_profile, and /etc/profile as correct locations for persistent environment variables. Therefore, the correct answer is D.

NEW QUESTION 5
Which of the following utilities supports the automation of security compliance and vulnerability management?

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

Answer: D

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

NEW QUESTION 6
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 7
A systems administrator needs to open the DNS TCP port on a Linux system from network 10.0.0.0/24. Which of the following commands should the administrator use for this task?

A. ufw allow dns/tcp to 10.0.0.0/24
B. ufw enable 53/tcp from 10.0.0.0/24
C. ufw allow 53/tcp from 10.0.0.0/24
D. ufw disable from 10.0.0.0/24

Answer: C

Explanation:
Firewall configuration is a key topic in theSecuritydomain of CompTIA Linux+ V8. DNS primarily uses UDP port 53, butTCP port 53is also required for zone transfers, large responses, and certain reliability scenarios. In this case, the administrator explicitly needs to allowDNS over TCPfrom a specific network.
The correct command is ufw allow 53/tcp from 10.0.0.0/24. This rule allows incoming TCP traffic on port 53 only from the specified subnet, following the principle ofleast privilege. Linux+ V8 documentation emphasizes restricting firewall rules by source network whenever possible to minimize attack surfaces.
OptionAis incorrect because UFW service aliases like dns are not always guaranteed to map explicitly to TCP, and the syntax is incomplete. OptionBis invalid because ufw enable is used to enable the firewall globally and does not define rules. OptionDdisables firewall protections and introduces a major security risk.
Linux+ V8 best practices stress precise, minimal firewall rules instead of broad or disabling actions. Therefore,Cis the correct and secure choice.

NEW QUESTION 8
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 9
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 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 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 /.

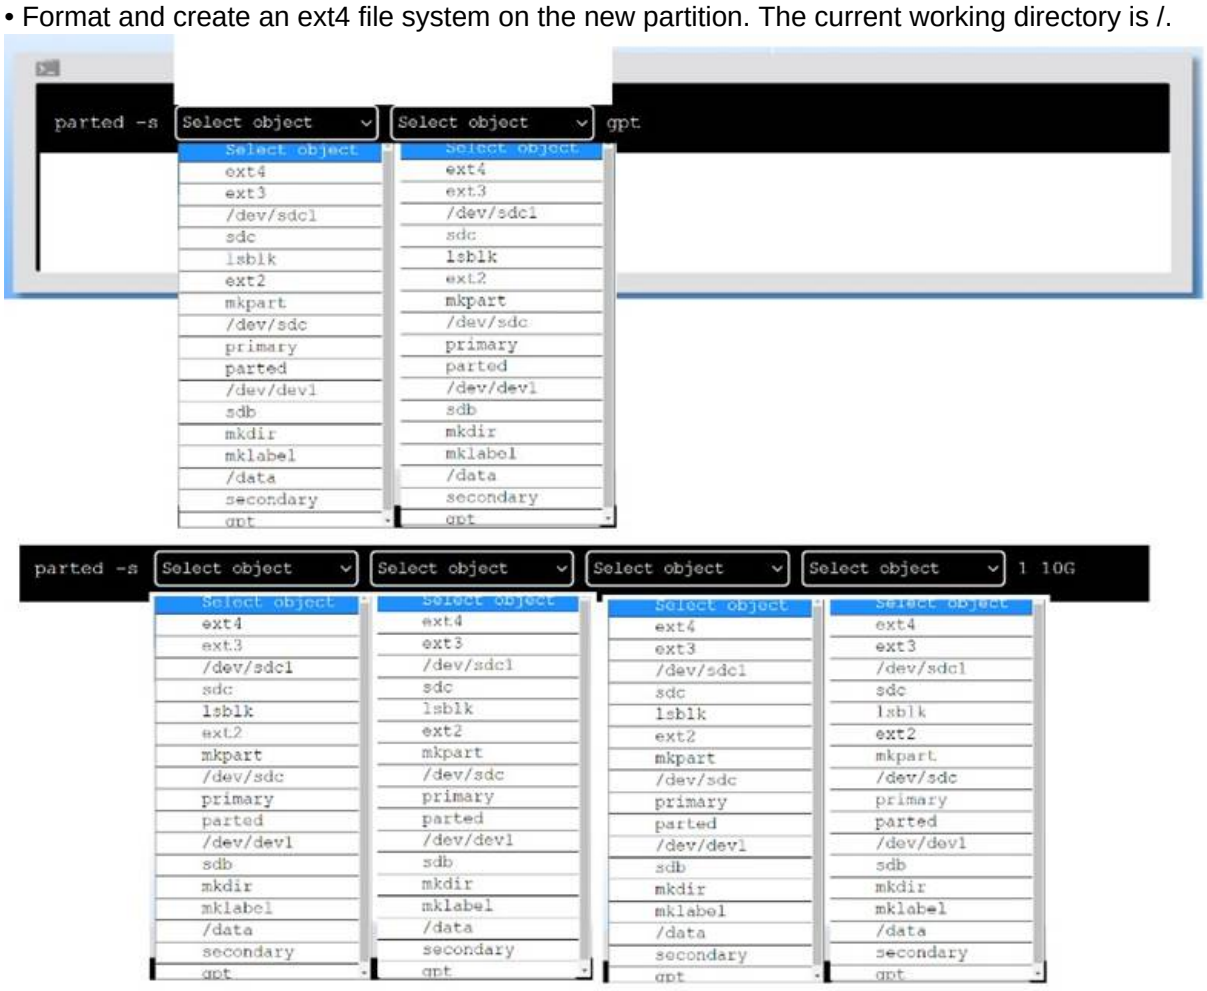

A. Mastered
B. Not Mastered

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 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 19
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 24
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
```

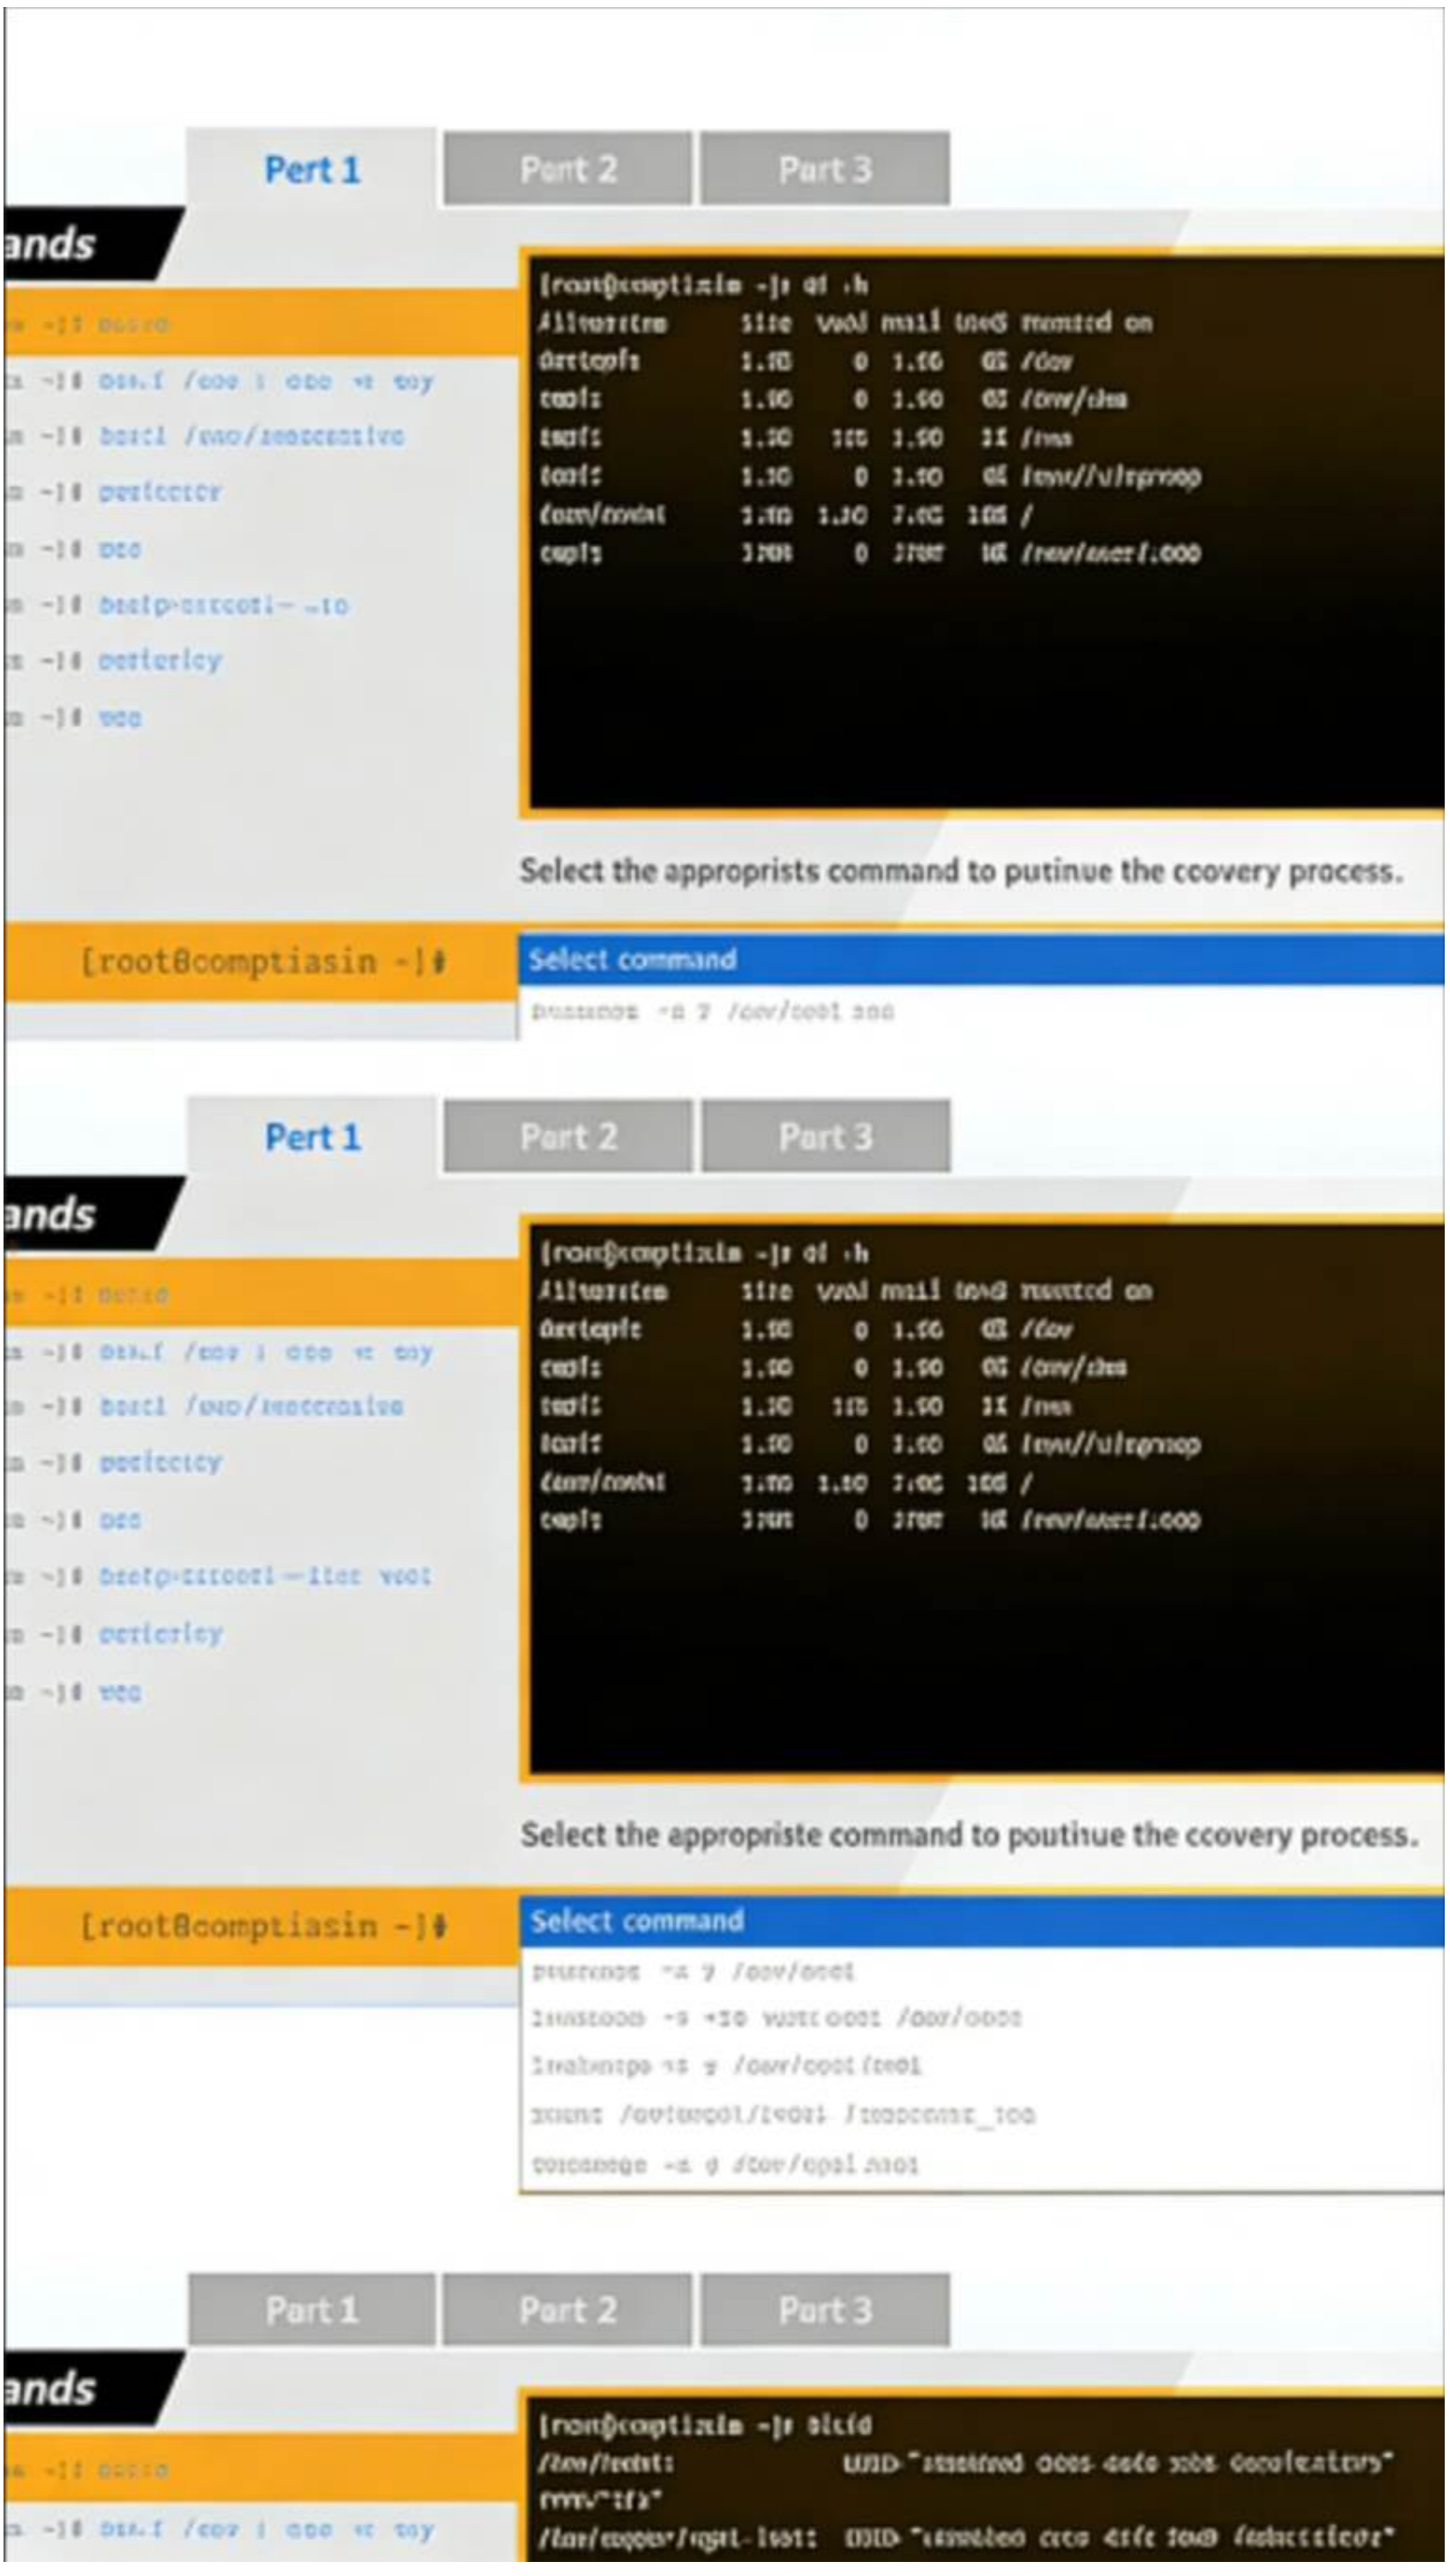
Part 1
Part 2
Part 3
Select the appropriste command to putinue the ccovery process.
[root@comptiasin ~]#
Select command
Part 1
Part 2
Part 3
Select the appropriste command to poutinue the ccovery process.
[root@comptiasin ~]#
Select command
Part 1
Part 2
Part 3

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 29
Users report that a Linux system is unresponsive and simple commands take too long to complete. The Linux administrator logs in to the system and sees the following:
Output 1:
10:06:29 up 235 day, 19:23, 2 users, load average: 8.71, 8.24, 7.71
Output 2:

Linux 6.8.0-31-generic (host) 05/10/2024 _x86_64_ (4 CPU)

10:07:42	AM	CPU	%usr	%nice	%sys	%iowait	%irq	%soft	%steal	%guest	%gnice	%idle
10:07:42	AM	all	65.88	0	20.54	5.65	0	7.93	0	0	0	0

Which of the following is the system experiencing?
A. High latency
B. High uptime
C. High CPU load
D. High I/O wait times

Answer: C

Explanation:
This scenario is a classic performance troubleshooting case covered under the Troubleshooting domain of the CompTIA Linux+ V8 objectives. The key indicators to analyze are the load average values and the CPU utilization statistics.
The uptime command shows load averages of 8.71, 8.24, and 7.71 over the 1-, 5-, and 15-minute intervals. Load average represents the average number of processes that are either running on the CPU or waiting to run. On a system with 4 CPU cores, a healthy load average would typically be close to or below 4. Load averages consistently near or above 8 indicate that there are significantly more runnable processes than available CPU resources, causing processes to wait and resulting in poor system responsiveness.
The CPU output further confirms this condition. The %idle value is 0, meaning the CPU has no idle time available. The majority of CPU time is spent in user space (65.88%) and system/kernel space (20.54%), indicating heavy computational and kernel activity. While %iowait is present at 5.65%, it is not high enough to suggest that disk I/O is the primary bottleneck.
Option C, high CPU load, best explains the symptoms. High CPU load causes commands to execute slowly because processes are competing for limited CPU time. This directly matches the observed behavior of the system being unresponsive.
The other options are incorrect. High uptime simply indicates how long the system has been running and does not cause performance issues by itself. High latency is a general term and not a specific diagnosis shown by the metrics provided. High I/O wait times would require a significantly higher %iowait value.
According to Linux+ V8 documentation, correlating load averages with CPU core count and utilization is essential for accurate performance diagnosis. Therefore, the correct answer is C. High CPU load.

NEW QUESTION 34
A Linux administrator just finished setting up passwordless SSH authentication between two nodes. However, upon test validation, the remote host prompts for a

password. Given the following logs:

```
-rw-------. 1 root root 588 Apr 3 2022 authorized_keys

avc: denied { read } for pid=xxxx comm="sshd" name="authorized_keys" dev="dm-5" ino=xxxx scontext=system_u:system_r:sshd_t:s0-s0:c0.c1
tcontext=unconfined_u:object_r:home_root_t:s0 tclass=file
[...]

SELinux status: enabled
SELinuxfs mount: /sys/fs/selinux
SELinux root directory: /etc/selinux
Loaded policy name: targeted
Current mode: enforcing
Mode from config file: enforcing
Policy MLS status: enabled
Policy deny_unknown status: allowed
Max kernel policy version: 31
```

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

A. The SELinux policy is incorrectly targeting the unconf ined_u context.
B. The administrator forgot to restart the SSHD after creating the authorizedjceys file.
C. The authorized_keys file has the incorrect root permissions assigned.
D. The authorized_keys file does not have the correct security context to match SELinux policy.

Answer: D

Explanation:
This issue is directly related to SELinux enforcement, which is a key topic in the Security domain of CompTIA Linux+ V8. The logs clearly indicate that SSH key-based authentication is failing due to an SELinux access control violation rather than a traditional file permission or SSH configuration problem.
The most important clue is the AVC denial message, which shows that the sshd process is being denied read access to the authorized_keys file. The security context of the file is listed as unconfined_u:object_r: home_root_t:s0. Under a targeted SELinux policy, SSH is only permitted to read authorized_keys files that are labeled with the correct SELinux type, typically ssh_home_t.
Because SELinux is running in enforcing mode, it actively blocks access that violates policy rules, even if standard UNIX permissions are correct. Although the file permissions (600) are acceptable for an authorized_keys file, SELinux does not rely solely on traditional permissions. The mismatch between the expected SELinux context and the actual context prevents sshd from accessing the file, causing SSH to fall back to password authentication.
Option D correctly identifies the root cause: the authorized_keys file does not have the correct SELinux security context. This is a well-documented Linux+ V8 troubleshooting scenario, commonly resolved by restoring the correct context using commands such as restorecon or by ensuring the file resides in a properly labeled home directory.
The other options are incorrect. Restarting sshd does not fix SELinux labeling issues. The policy itself is functioning as intended, and file ownership alone does not override SELinux access controls.
Linux+ V8 documentation emphasizes that SELinux denials must be addressed by correcting file contexts rather than weakening security controls. Therefore, the correct answer is D.

NEW QUESTION 36
An administrator attempts to install updates on a Linux system but receives error messages regarding a specific repository. Which of the following commands should the administrator use to verify that the repository is installed and enabled?

A. yum repo-pkgs
B. yum list installed repos
C. yum reposync available
D. yum repolist all

Answer: D

Explanation:
Package management troubleshooting is an important skill in Linux+ V8, especially on RPM-based distributions that use yum or dnf. When update errors reference a repository, the administrator must verify whether the repository exists and whether it is enabled.
The command yum repolist all displays all configured repositories, including those that are enabled, disabled, or temporarily unavailable. This makes it the most effective command for diagnosing repository-related issues. It allows administrators to quickly confirm the repository??s status and take corrective action, such as enabling it or fixing configuration errors.
The other options are incorrect. yum repo-pkgs manages packages within a repository but does not list repository status. yum list installed repos is not a valid yum command. yum reposync is used to mirror repositories locally and is not intended for verification.
Linux+ V8 documentation highlights yum repolist all as the standard command for repository inspection and troubleshooting.
Therefore, the correct answer is D. yum repolist all.

NEW QUESTION 40
A Linux administrator attempts to log in to a server over SSH as root and receives the following error message: Permission denied, please try again. The administrator is able to log in to the console of the server directly with root and confirms the password is correct. The administrator reviews the configuration of the SSH service and gets the following output:

```
Port 22
PermitRootLogin prohibit-password
PasswordAuthentication yes
PermitEmptyPassword no
Use PAM no
MaxSessions 1
MaxAuthTries 3
```

Based on the above output, which of the following will most likely allow the administrator to log in over SSH to the server?

A. Log out other user sessions because only one is allowed at a time.
B. Enable PAM and configure the SSH module.
C. Modify the SSH port to use 2222.
D. Use a key to log in as root over SSH.

Answer: D

Explanation:
The SSH configuration option PermitRootLogin prohibit-password prevents the root user from logging in with password authentication. This setting means root cannot use a password to log in via SSH; only key- based authentication is permitted for root. The administrator can still log in as root locally, which is not affected by this SSH configuration. To allow SSH access as root, the administrator must use an SSH key instead of a password.
Other options:
* A. MaxSessions controls the number of simultaneous SSH sessions but is not causing the login denial here.
* B. PAM (Pluggable Authentication Modules) is disabled, but enabling it is not required for basic SSH authentication.
* C. Changing the SSH port is unrelated to the authentication method issue.
Reference:
CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 11: "Securing Linux", Section: "Securing SSH Access"
CompTIA Linux+ XK0-006 Objectives, Domain 3.0: Security

NEW QUESTION 45
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 48
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 51
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 52
Which of the following can reduce the attack surface area in relation to Linux hardening?

A. Customizing the log-in banner
B. Reducing the number of directories created
C. Extending the SSH startup timeout period
D. Enforcing password strength and complexity

Answer: D

Explanation:
Comprehensive and Detailed Explanation From Exact Extract:
Reducing the attack surface area in Linux hardening refers to limiting possible points of unauthorized access. According to the CompTIA Linux+ Official Study Guide (Exam XK0-006), enforcing strong password policies is a critical aspect of security hardening. This practice ensures that user accounts are protected by passwords that are difficult to guess or crack, thus minimizing the risk of successful brute-force attacks. Implementing password complexity requirements (such as minimum length, use of uppercase, lowercase, numbers, and special characters) directly addresses one of the primary vectors for unauthorized access.
Other options do not have a direct impact on reducing the attack surface:
* A. Customizing the log-in bannerserves as a legal notification and does not affect system vulnerabilities.
* B. Reducing the number of directories createdis not related to hardening or access control.
* C. Extending the SSH startup timeout periodmay give attackers more time to attempt a connection and does not increase security.
[Reference:, CompTIA Linux+ Study Guide: Exam XK0-006, Sybex, Chapter 11: "Securing the System", Section: "Implementing Password Policies", CompTIA Linux+ XK0-006 Exam Objectives, Domain 3.0: Security, , ,]

NEW QUESTION 55
A Linux systems administrator needs to extract the contents of a file named /home/dev/web.bkp to the /var/www/html/ directory. Which of the following commands should the administrator use?

A. cd /var/www/html/ && gzip -c /home/dev/web.bkp | tar xf -
B. pushd /var/www/html/ && cpio -idv < /home/dev/web.bkp && popd
C. tar -c -f /home/dev/web.bkp /var/www/html/
D. unzip -c /home/dev/web.bkp /var/www/html/

Answer: B

Explanation:
Comprehensive and Detailed 250 to 350 words of Explanation From Linux+ V8 documents:
File extraction and backup restoration are fundamentalSystem Managementtasks covered in CompTIA Linux+ V8. In this scenario, the administrator must extract the contents of an existing backup file into a target directory.
The correct command isoption B, which uses cpio in extract mode. The command changes into the destination directory (/var/www/html/) using pushd, extracts the archive contents with cpio -idv, and then returns to the original directory with popd. This ensures that files are restored into the correct location without modifying paths inside the archive.
The cpio utility is commonly used for backups created with cpio -o and supports reading archive data from standard input. Linux+ V8 documentation includes cpio as a valid and supported archive format for backup and restore operations.
The other options are incorrect. OptionAincorrectly assumes the backup is a gzip-compressed tar archive. OptionCcreates a new archive instead of extracting one. OptionDassumes the file is a ZIP archive, which is not indicated by the .bkp extension.
Linux+ V8 emphasizes using the correct tool based on the archive format and restoring files into the intended directory. Therefore, the correct answer isB.

NEW QUESTION 59
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 64
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 67
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 68
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 70
A systems administrator needs to check the statuses of all the services on a Linux server. Which of the following commands accomplishes this task?

A. systemctl is-active --services
B. systemctl list-sockets --type=services
C. systemctl is-enabled --services
D. systemctl list-units --type=services

Answer: D

Explanation:
Service management usingsystemdis a core Linux+ V8 system management objective. Administrators frequently need to view the current status of all services to determine which ones are running, stopped, failed, or inactive.
The correct command issystemctl list-units --type=services, which displays all loaded service units along with their current state, including whether they are active, inactive, failed, or running. This provides a comprehensive, real-time view of service statuses on the system and is commonly used during troubleshooting and audits.
OptionA, systemctl is-active, is designed to check the status of a single service, not all services. OptionBlists socket units, not services. OptionC, systemctl is-enabled, checks whether services are enabled at boot, not whether they are currently running.
Linux+ V8 documentation explicitly references systemctl list-units --type=service as the primary command for viewing service runtime states. Therefore, the correct answer isD.

NEW QUESTION 71
An administrator logs in to a Linux server and notices the clock is 37 minutes fast. Which of the following commands will fix the issue?

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

Answer: B

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

NEW QUESTION 75
A Linux software developer wants to use AI to optimize source code used in a commercial product. Which of the following steps should the developer take first?

A. Research which available AI chatbots are best at optimizing source code.
B. Verify that the company has a policy governing the use of AI in software development.
C. Install a private LLM to use on the internal network for source code optimization.
D. Use open-source LLMs that undergo regular security reviews by the community.

Answer: B

Explanation:
Linux+ V8 emphasizessecurity, compliance, and governancewhen introducing new automation technologies, including AI. Before using AI tools to optimize commercial source code, the developer must ensure that such usage complies with organizational policies.
OptionBis correct because verifying company policy is the first and most critical step. AI tools may introduce risks such as intellectual property leakage, licensing conflicts, or regulatory violations. Many organizations restrict how source code can be shared with external systems, including AI services.
The other options are premature. Selecting tools or deploying models should only occur after policy approval. Linux+ V8 highlights governance-first approaches when adopting automation technologies.
Therefore, the correct answer isB.

NEW QUESTION 80
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 84
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 87
Which of the following best describes the role of initrd?

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

Answer: B

NEW QUESTION 89
.....

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