XK0-005$^{\text{Q\&As}}$

CompTIA Linux+ Certification Exam

Pass CompTIA XK0-005 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4itsure.com/xk0-005.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by CompTIA Official Exam Center

QUESTION 1

A systems administrator is tasked with setting up key-based SSH authentication. In which of the following locations should the administrator place the public keys for the server?

A. ~/.sshd/authkeys

B. ~/.ssh/keys

C. ~/.ssh/authorized_keys

D. ~/.ssh/keyauth

Correct Answer: C

Explanation: The administrator should place the public keys for the server in the ~/.ssh/authorized_keys file. The SSH (Secure Shell) protocol is a method for establishing secure and encrypted connections between remote systems. The SSH protocol supports two types of authentication: password-based and key-based. Password-based authentication requires the user to enter the password of the remote system every time they connect. Key-based authentication requires the user to generate a pair of cryptographic keys: a public key and a private key. The public key is stored on the remote system, while the private key is kept on the local system. The public key and the private key are mathematically related, but not identical. The SSH protocol uses the keys to verify the identity of the user and establish a secure connection without requiring a password. The ~/.ssh/authorized_keys file is a file that contains the public keys of the users who are allowed to connect to the remote system using key-based authentication. The administrator should place the public keys for the server in this file, one per line, and set the appropriate permissions for the file. The administrator should also configure the SSH server to enable key-based authentication by editing the /etc/ssh/sshd_config file and setting the option PasswordAuthentication to no. The administrator should place the public keys for the server in the ~/.ssh/authorized_keys file. This is the correct answer to the question. The other options are incorrect because they are not the standard locations for the public keys for the server (~/.sshd/authkeys, ~/.ssh/keys, or ~/.ssh/keyauth). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Implementing Basic Security, page 513.

QUESTION 2

Users are experiencing high latency when accessing a web application served by a Linux machine. A systems administrator checks the network interface counters and sees the following:

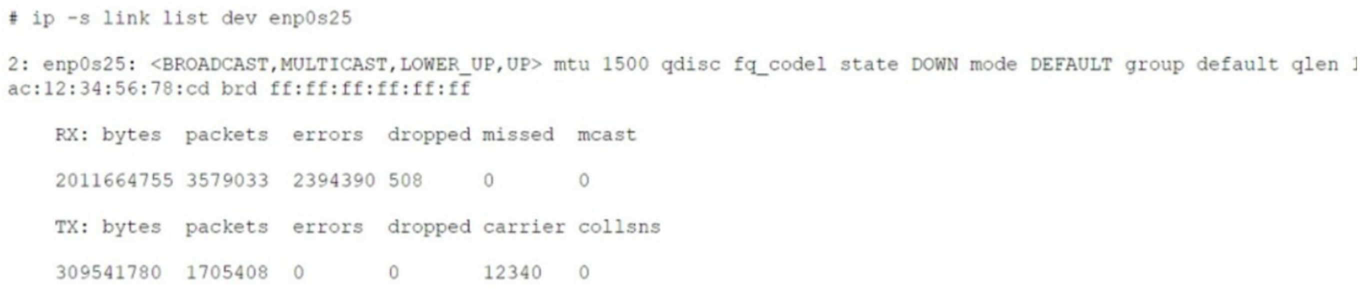

```
# ip -s link list dev enp0s25

2: enp0s25: <BROADCAST,MULTICAST,LOWER_UP,UP> mtu 1500 qdisc fq_codel state DOWN mode DEFAULT group default qlen 1
ac:12:34:56:78:cd brd ff:ff:ff:ff:ff:ff

    RX: bytes  packets  errors  dropped missed  mcast

    2011664755 3579033  2394390 508     0       0

    TX: bytes  packets  errors  dropped carrier collsns

    309541780  1705408  0       0       12340   0
```

Which of the following is the most probable cause of the observed latency?

A. The network interface is disconnected.

B. A connection problem exists on the network interface.

C. No IP address is assigned to the interface.

D. The gateway is unreachable.

Correct Answer: B

The high number of errors and dropped packets in the output of the network interface counters indicate a connection problem on the network interface.

References:

CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Networking, Section: Troubleshooting Network Issues, Page 359. Linux+ (Plus) Certification, Exam Objectives: 4.3 Given a scenario, troubleshoot and resolve

basic network configuration and connectivity issues.

QUESTION 3

A systems administrator has been unable to terminate a process. Which of the following should the administrator use to forcibly stop the process?

A. kill -1

B. kill -3

C. kill -15

D. kill -HUP

E. kill -TERM

Correct Answer: E

Explanation: The administrator should use the command kill -TERM to forcibly stop the process. The kill command is a tool for sending signals to processes on Linux systems. Signals are messages that inform the processes about certain events and actions. The processes can react to the signals by performing predefined or user-defined actions, such as terminating, suspending, resuming, or ignoring. The -TERM option specifies the signal name or number that the kill command should send. The TERM signal, which stands for terminate, is the default signal that the kill command sends if no option is specified. The TERM signal requests the process to terminate gracefully, by closing any open files, releasing any resources, and performing any cleanup tasks. However, if the process does not respond to the TERM signal, the kill command can send a stronger signal, such as the KILL signal, which forces the process to terminate immediately, without any cleanup. The administrator should use the command kill -TERM to forcibly stop the process. This is the correct answer to the question. The other options are incorrect because they either do not terminate the process (kill -1 or kill 3) or do not terminate the process forcibly (kill -15 or kill -HUP). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes, page 431.

QUESTION 4

A user is attempting to log in to a Linux server that has Kerberos SSO ena-bled. Which of the following commands should the user run to authenticate and then show the ticket grants? (Select TWO).

A. kinit

B. klist

C. kexec

D. kioad

E. pkexec

F. realm

Correct Answer: AB

The following commands can help the user to authenticate and show the ticket grants using Kerberos SSO on a Linux server:

kinit: This command obtains and caches an initial ticket-granting ticket (TGT) for the user from the Kerberos key distribution center (KDC). The user needs to enter their password or use a keytab file to authenticate1. klist: This command lists

the cached tickets, including the TGT and any service tickets, for the user. It also shows the expiration time and flags for each ticket2. For example, the user can run the following commands to log in and view their tickets:

```
$ kinit username@REALM

Password for username@REALM:

$ klist

Ticket cache: FILE:/tmp/krb5cc_1000

Default principal: username@REALM

Valid starting Expires Service principal

04/06/2023 16:06:59 04/07/2023 02:06:59 krbtgt/REALM@REALM renew until 04/13/2023 16:06:59
```

References:

kinit(1) - Linux man page, section "Description".

klist(1) - Linux man page, section "Description".

QUESTION 5

A Linux administrator is troubleshooting an issue in which an application service failed to start on a Linux server. The administrator runs a few commands and gets the following outputs:

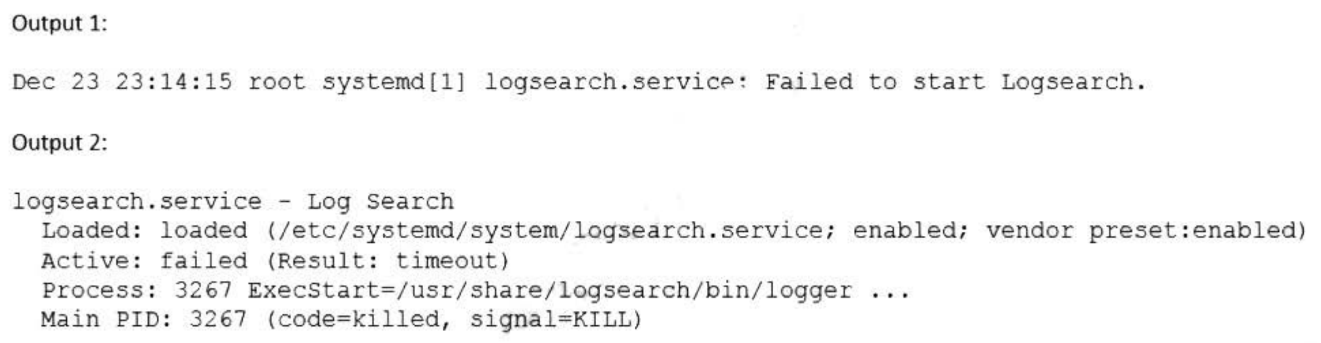

Output 1:

```
Dec 23 23:14:15 root systemd[1] logsearch.service: Failed to start Logsearch.
```

Output 2:

```
logsearch.service - Log Search
  Loaded: loaded (/etc/systemd/system/logsearch.service; enabled; vendor preset:enabled)
  Active: failed (Result: timeout)
  Process: 3267 ExecStart=/usr/share/logsearch/bin/logger ...
  Main PID: 3267 (code=killed, signal=KILL)
```

Based on the above outputs, which of the following is the MOST likely action the administrator should take to resolve this issue?

A. Enable the logsearch.service and restart the service.

B. Increase the TimeoutStartUSec configuration for the logsearch.sevice.

C. Update the OnCalendar configuration to schedule the start of the logsearch.service.

D. Update the KillSignal configuration for the logsearch.service to use TERM.

Correct Answer: B

Explanation: The administrator should increase the TimeoutStartUSec configuration for the logsearch.service to resolve the issue. The output of systemct1 status logsearch.service shows that the service failed to start due to a timeout. The output of cat /etc/systemd/system/logsearch.service shows that the service has a TimeoutStartUSec configuration of 10 seconds, which might be too short for the service to start. The administrator should increase this value to a higher number, such as 30 seconds or 1 minute, and then restart the service. The other options are incorrect because they are not related to the issue. The service is already enabled, as shown by the output of systemct1 is-enabled logsearch.service. The service does not use an OnCalendar configuration, as it is not a timer unit. The service does not use a KillSignal configuration, as it is not being killed by a signal. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, pages 434-435.

QUESTION 6

A Linux administrator has physically added a new RAID adapter to a system. Which of the following commands should the Linux administrator run to confirm that the device has been recognized? (Select TWO).

A. rmmod

B. Is -11 /etc

C. Ishw --class disk

D. pvdisplay

E. rmdir /dev

F. dmesg

Correct Answer: CF

The following commands can help you confirm that the new RAID adapter has been recognized by the Linux system:

dmesg: This command displays the kernel messages, which can show the information about the newly detected hardware device. You can use dmesg | grep -i raid to filter the output for RAID-related messages. lshw -class disk: This

command lists the disk devices on the system, including the RAID controller and its model name. You can use lshw -class disk | grep -i raid to filter the output for RAID-related information1.

The other commands are not relevant for this purpose. For example:

rmmod: This command removes a module from the Linux kernel, which is not useful for detecting a new device.

ls -l /etc: This command lists the files and directories in the /etc directory, which is not related to hardware devices.

pvdisplay: This command displays the attributes of physical volumes, which are part of the logical volume management (LVM) system, not the RAID system. rmdir /dev: This command removes an empty directory, which is not helpful for

detecting a new device. Moreover, /dev is a special directory that contains device files, and should not be removed.

QUESTION 7

An administrator attempts to rename a file on a server but receives the following error.

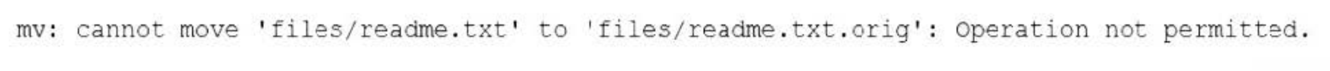

```
mv: cannot move 'files/readme.txt' to 'files/readme.txt.orig': Operation not permitted.
```

The administrator then runs a few commands and obtains the following output:

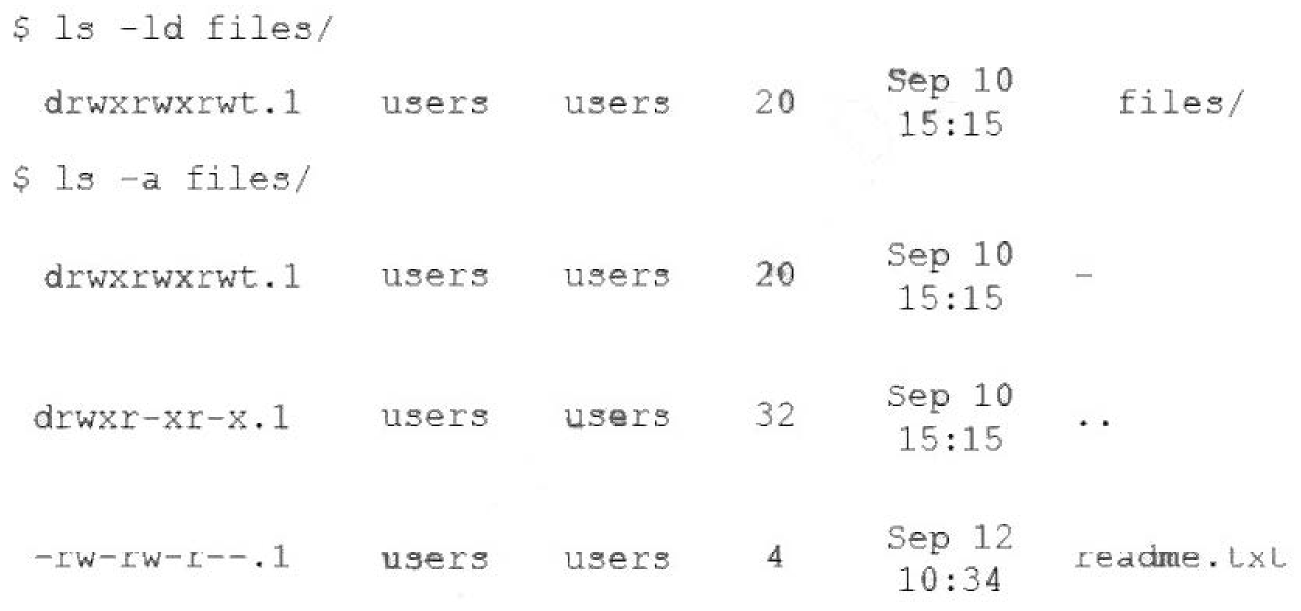

```
$ ls -ld files/
 drwxrwxrwt.1   users   users   20   Sep 10     files/
                                     15:15
$ ls -a files/
 drwxrwxrwt.1   users   users   20   Sep 10   .
                                     15:15
 drwxr-xr-x.1   users   users   32   Sep 10   ..
                                     15:15
 -rw-rw-r--.1   users   users   4    Sep 12   readme.txt
                                     10:34
```

Which of the following commands should the administrator run NEXT to allow the file to be renamed by any user?

A. chgrp reet files

B. chacl -R 644 files

C. chown users files

D. chmod -t files

Correct Answer: D

Explanation: The command that the administrator should run NEXT to allow the file to be renamed by any user is chmod -t files. This command uses the chmod tool, which is used to change file permissions and access modes. The -t option removes (or sets) the sticky bit on a directory, which restricts deletion or renaming of files within that directory to only their owners or root. In this case, since files is a directory with sticky bit set (indicated by t in drwxrwxrwt), removing it will allow any user to rename or delete files within that directory. The other options are not correct commands for allowing any user to rename files within files directory. The chgrp reet files command will change the group ownership of files directory to reet, but it will not affect its permissions or access modes. The chacl -R 644 files command is invalid, as chacl is used to change file access control lists (ACLs), not permissions or access modes. The chown users files command will change the user ownership of files directory to users, but it will not affect its permissions or access modes. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing Users and Groups; chmod(1) - Linux manual page

QUESTION 8

Developers have requested implementation of a persistent, static route on the application server. Packets sent over the interface eth0 to 10.0.213.5/32 should be routed via 10.0.5.1. Which of the following commands should the administrator run to achieve this goal?

A. route -i etho -p add 10.0.213.5 10.0.5.1

B. route modify eth0 +ipv4.routes "10.0.213.5/32 10.0.5.1"

C. echo "10.0.213.5 10.0.5.1 eth0" > /proc/net/route

D. ip route add 10.0.213.5/32 via 10.0.5.1 dev eth0

Correct Answer: D

Explanation: The command ip route add 10.0.213.5/32 via 10.0.5.1 dev eth0 adds a static route to the routing table that sends packets destined for 10.0.213.5/32 (a single host) through the gateway 10.0.5.1 on the interface eth0. This is the correct way to achieve the goal. The other options are incorrect because they either use the wrong syntax (route -i etho -p add), the wrong command (route modify), or the wrong file (/proc/net/route). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 379.

QUESTION 9

A Linux administrator was notified that a virtual server has an I/O bottleneck. The Linux administrator analyzes the following output:

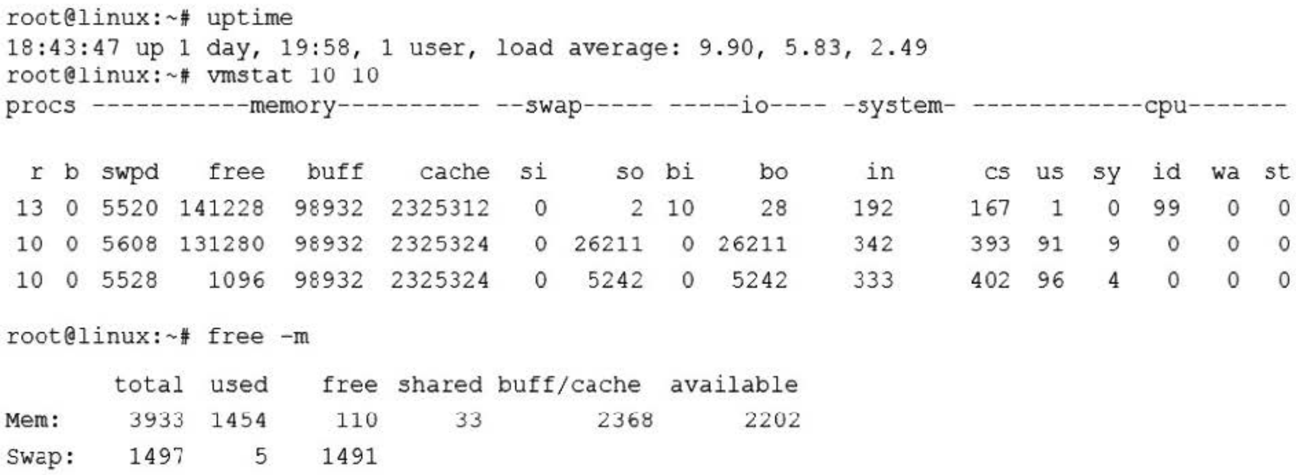

```
root@linux:~# uptime
18:43:47 up 1 day, 19:58, 1 user, load average: 9.90, 5.83, 2.49
root@linux:~# vmstat 10 10
procs -----------memory---------- --swap----- -----io---- -system- ------------cpu-------

  r  b  swpd   free    buff    cache   si     so  bi     bo     in      cs  us  sy  id  wa  st
 13  0  5520 141228  98932  2325312   0      2  10     28    192     167   1   0  99   0   0
 10  0  5608 131280  98932  2325324   0  26211   0  26211    342     393  91   9   0   0   0
 10  0  5528   1096  98932  2325324   0   5242   0   5242    333     402  96   4   0   0   0

root@linux:~# free -m
          total  used    free  shared buff/cache  available
Mem:       3933  1454     110      33       2368       2202
Swap:      1497     5    1491
```

Given there is a single CPU in the sever, which of the following is causing the slowness?

A. The system is running out of swap space.

B. The CPU is overloaded.

C. The memory is exhausted.

D. The processes are paging.

Correct Answer: B

Explanation: The slowness is caused by the CPU being overloaded. The iostat command shows that the CPU utilization is 100%, which means that there are more processes competing for CPU time than the CPU can handle. The other

options are incorrect because:

The system is not running out of swap space, as shown by the iostat command, which shows that there is no swap activity (si and so columns are zero). The memory is not exhausted, as shown by the free -m command, which shows that

there is still available memory (avail column) and free buffer/cache memory (buff/cache column).

The processes are not paging, as shown by the vmstat command, which shows that there are no major page faults (majflt column) and no swap activity (si and so columns). References: CompTIA Linux+ Study Guide, Fourth Edition, page

417- 419, 424-425.

QUESTION 10

A Linux administrator has defined a systemd script docker-repository.mount to mount a volume for use by the Docker service. The administrator wants to ensure that Docker service does not start until the volume is mounted. Which of the following configurations needs to be added to the Docker service definition to best accomplish this task?

A. After=docker-respository.mount

B. ExecStart=/usr/bin/mount -a

C. Requires=docker-repository.mount

D. RequiresMountsFor=docker-repository.mount

Correct Answer: C

This option declares an explicit dependency between the Docker service and the docker- repository.mount unit. It means that the Docker service will not start unless the docker- repository.mount unit is successfully activated. This ensures that the volume is mounted before the Docker service tries to use it12.

References: 1: systemd.unit - systemd unit configuration 2: How to mount host volumes into docker containers in Dockerfile during build

QUESTION 11

A systems administrator wants to test the route between IP address 10.0.2.15 and IP address 192.168.1.40. Which of the following commands will accomplish this task?

A. route -e get to 192.168.1.40 from 10.0.2.15

B. ip route get 192.163.1.40 from 10.0.2.15

C. ip route 192.169.1.40 to 10.0.2.15

D. route -n 192.168.1.40 from 10.0.2.15

Correct Answer: B

Explanation: The command ip route get 192.168.1.40 from 10.0.2.15 will test the route between the IP address 10.0.2.15 and the IP address 192.168.1.40. The ip route get command shows the routing decision for a given destination and source address. This is the correct command to accomplish the task. The other options are incorrect because they either use the wrong commands (route instead of ip route), the wrong options (-e or - n instead of get), or the wrong syntax (to instead of from). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 379.

QUESTION 12

A Linux administrator needs to remove software from the server. Which of the following RPM options should be used?

A. rpm -s

B. rm -d

C. rpm -q

D. rpm -e

Correct Answer: D

Explanation: The RPM option -e should be used to remove software from the server. The rpm command is a tool for managing software packages on RPM-based Linux distributions. The -e option stands for erase and removes the specified package from the system. This is the correct option to use to accomplish the task. The other options are incorrect because they either do not exist (-s or -d) or do not remove software (-q stands for query and displays information about the package). References: CompTIA Linux+ (XK0- 005) Certification Study Guide, Chapter 16: Managing Software, page 489.

QUESTION 13

An administrator would like to securely connect to a server and forward port 8080 on a local machine to port 80 on the server. Which of the following commands should the administrator use to satisfy both requirements?

A. ssh --L 8080: localhost:80 admin@server

B. ssh --R 8080: localhost:80 admin@server

C. ssh --L 80 : localhost:8080 admin@server

D. ssh --R 80 : localhost:8080 admin@server

Correct Answer: A

This command will create a local port forwarding, which means that connections from the SSH client are forwarded via the SSH server, then to a destination server. In this case, the destination server is the same as the SSH server (localhost),

and the destination port is 80. The SSH client will listen on port 8080 on the local machine, and any connection to that port will be forwarded to port 80 on the server. This way, the administrator can securely access the web service running on

port 80 on the server by using http://localhost:8080 on the local machine.

The other options are incorrect because:

B. ssh -R 8080:localhost:80 admin@server

This command will create a remote port forwarding, which means that connections from the SSH server are forwarded via the SSH client, then to a destination server. In this case, the destination server is the same as the SSH client

(localhost), and the destination port is 80. The SSH server will listen on port 8080 on the remote machine, and any connection to that port will be forwarded to port 80 on the client. This is not what the administrator wants to do.

C. ssh -L 80:localhost:8080 admin@server

This command will also create a local port forwarding, but it will use port 80 on the local machine and port 8080 on the server. This is not what the administrator wants to do, and it may also fail if port 80 is already in use by another service on

the local machine.

D. ssh -R admin@server

This command is incomplete and invalid. It does not specify any port numbers or destination addresses for the remote port forwarding. It will also fail if the SSH server does not allow remote port forwarding.

References:

CompTIA Linux+ Certification Exam Objectives

How to Set up SSH Tunneling (Port Forwarding)

QUESTION 14

A Linux systems administrator receives a notification that one of the server\\'s filesystems is full. Which of the following commands would help the administrator to identify this filesystem?

A. lsblk

B. fdisk

C. df -h

D. du -ah

Correct Answer: C

Explanation: The df -h command can be used to identify the filesystem that is full. This command displays the disk usage of each mounted filesystem in a human-readable format, showing the total size, used space, available space, and percentage of each filesystem. The lsblk command displays information about block devices, not filesystems. The fdisk command can be used to manipulate partition tables, not check disk usage. The du -ah command displays the disk usage of each file and directory in a human-readable format, not the filesystems. References: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 14: Managing Disk Storage, page 454.

QUESTION 15

A Linux system is failing to start due to issues with several critical system processes. Which of the following options can be used to boot the system into the single user mode? (Choose two.)

A. Execute the following command from the GRUB rescue shell: mount -o remount, ro/sysroot.

B. Interrupt the boot process in the GRUB menu and add systemd.unit=single in the kernel line.

C. Interrupt the boot process in the GRUB menu and add systemd.unit=rescue.target in the kernel line.

D. Interrupt the boot process in the GRUB menu and add single=user in the kernel line.

E. Interrupt the boot process in the GRUB menu and add init=/bin/bash in the kernel line.

F. Interrupt the boot process in the GRUB menu and add systemd.unit=single.target in the kernel line.

Correct Answer: CF

The administrator can use the following two options to boot the system into the single user mode: Interrupt the boot process in the GRUB menu and add systemd.unit=rescue.target in the kernel line. This option will boot the system into the rescue mode, which is a minimal environment that allows the administrator to perform basic tasks such as repairing the system. The GRUB menu is a screen that appears when the system is powered on and allows the administrator to choose which kernel or operating system to boot. The kernel line is a line that specifies the parameters for the kernel, such as the root device, the init system, and the boot options. The administrator can interrupt the boot process by pressing the e key in the GRUB menu and edit the kernel line by adding systemd.unit=rescue.target at the end. This option will tell the system to use the rescue target, which is a unit that defines the state of the system in the rescue mode. The administrator can then press Ctrl+X to boot the system with the modified kernel line. This option will boot the system into the single user mode and allow the administrator to troubleshoot the issues. Interrupt the boot process in the GRUB menu and add systemd.unit=single.target in the kernel line. This option will boot the system into the single user mode, which is a mode that allows the administrator to log in as the root user and perform maintenance tasks. The GRUB menu and the kernel line are the same as the previous option. The administrator can interrupt the boot process by pressing the e key in the GRUB menu and edit the kernel line by adding systemd.unit=single.target at the end. This option will tell the system to use the single target, which is a unit that defines the state of the system in the single user mode. The administrator can then press Ctrl+X to boot the system with the modified kernel line. This option will boot the system into the single user mode and allow the administrator to troubleshoot the issues. The other options are incorrect because they either do not boot the system into the single user mode (execute the following command from the GRUB rescue shell: mount -o remount, ro/sysroot or interrupt the boot process in the GRUB menu and add systemd.unit=single in the kernel line) or do not use the correct syntax (interrupt the boot process in the GRUB menu and add single=user in the kernel line or interrupt the boot process in the GRUB menu and add init=/bin/bash in the kernel line). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing the Linux Boot Process, pages 267-268.

XK0-005 PDF Dumps XK0-005 VCE Dumps XK0-005 Practice Test