



XK0-005^{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.passapply.com/xk0-005.html>

100% Passing Guarantee
100% Money Back Assurance

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

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers





QUESTION 1

A Linux administrator is troubleshooting a memory-related issue. Based on the output of the commands:

```
$ vmstat -s --unit M

968 M total memory
331 M used memory
482 M active memory
279 M inactive memory
99 M free memory

$ free -h

              total        used        free      shared  buff/cache   available
Mem:          968M        331M        95M         13M         540M         458M
Swap:           0           0           0

$ ps -aux | grep script.sh
USER      PID     %CPU   %MEM    VSZ       RSS        TTY  STAT   START   TIME    COMMAND
user      8321   2.8    40.5  3224846  371687    7      SN    16:49   2:09   /home/user/script.sh
```

Which of the following commands would address the issue?

- A. top -p 8321
- B. kill -9 8321
- C. renice -10 8321
- D. free 8321

Correct Answer: B

Explanation: The command that would address the memory-related issue is kill -9 8321. This command will send a SIGKILL signal to the process with the PID 8321, which is the mysql process that is using 99.7% of the available memory according to the top output. The SIGKILL signal will terminate the process immediately and free up the memory it was using. However, this command should be used with caution as it may cause data loss or corruption if the process was performing some critical operations. The other options are not correct commands for addressing the memory-related issue. The top -p 8321 command will only display information about the process with the PID 8321, but will not kill it or reduce its memory usage. The renice -10 8321 command will change the priority (niceness) of the process with the PID 8321 to -10, which means it will have a higher scheduling priority, but this will not affect its memory consumption. The free 8321 command is invalid because free does not take a PID as an argument; free only displays information about the total, used, and free memory in the system. References: How to troubleshoot Linux server memory issues; kill(1) - Linux manual page

QUESTION 2

A systems administrator is checking the system logs. The administrator wants to look at the last 20 lines of a log. Which of the following will execute the command?



A. tail -v 20

B. tail -n 20

C. tail -c 20

D. tail -l 20

Correct Answer: B

Explanation: The command tail -n 20 will display the last 20 lines of a file. The -n option specifies the number of lines to show. This is the correct command to execute the task. The other options are incorrect because they either use the wrong options (-v, -c, or -l) or have the wrong arguments (20 instead of 20 filename). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, page 352.

QUESTION 3

A systems administrator is compiling a report containing information about processes that are listening on the network ports of a Linux server. Which of the following commands will allow the administrator to obtain the needed information?

A. ss -pint

B. tcpdump -nL

C. netstat -pn

D. lsof -lt

Correct Answer: A

Explanation: The command ss -pint will allow the administrator to obtain the needed information about processes that are listening on the network ports of a Linux server. The ss command is a tool for displaying socket statistics on Linux systems. Sockets are endpoints of network communication that allow processes to exchange data over the network. The ss command can show various information about the sockets, such as the state, address, port, protocol, and process. The -pint option specifies the filters and flags that the ss command should apply. The -p option shows the process name and ID that owns the socket. The -i option shows the internal information about the socket, such as the send and receive queue, the congestion window, and the retransmission timeout. The -n option shows the numerical address and port, instead of resolving the hostnames and service names. The -t option shows only the TCP sockets, which are the most common type of sockets used for network communication. The command ss -pint will display the socket statistics for the TCP sockets, along with the process name and ID, the numerical address and port, and the internal information. This will allow the administrator to obtain the needed information about processes that are listening on the network ports of a Linux server. This is the correct command to use to obtain the needed information. The other options are incorrect because they either do not show the socket statistics (tcpdump -nL or lsof -lt) or do not show the process name and ID (netstat -pn). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 389.

QUESTION 4

A cloud engineer needs to check the link status of a network interface named eth1 in a Linux server. Which of the following commands can help to achieve the goal?

A. ifconfig hw eth1



B. netstat -r eth1

C. ss -ti eth1

D. ip link show eth1

Correct Answer: D

Explanation: The ip link show eth1 command can be used to check the link status of a network interface named eth1 in a Linux server. It will display information such as the MAC address, MTU, state, and flags of the interface. The ifconfig hw eth1 command is invalid, as hw is not a valid option for ifconfig. The netstat -r eth1 command would display the routing table for eth1, not the link status. The ss -ti eth1 command would display TCP information for sockets associated with eth1, not the link status. References: CompTIA Linux+ (XK0- 005) Certification Study Guide, Chapter 13: Networking Fundamentals, page 436.

QUESTION 5

A Linux administrator is scheduling a system job that runs a script to check available disk space every hour. The Linux administrator does not want users to be able to start the job. Given the following:

```
[Unit]
Description=Check available disk space
RefuseManualstart=yes
RefuseManualStop=yes

[Timer]
Persistent=true
OnCalendar=*-*-* *:00:00
Unit=checkdiskspace.service

[Install]
WantedBy=timers.target
```

The Linux administrator attempts to start the timer service but receives the following error message:

```
Failed to start checkdiskspace.timer: Operation refused ...
```



Which of the following is MOST likely the reason the timer will not start?

- A. The checkdiskspace.timer unit should be enabled via systemctl.
- B. The timers.target should be reloaded to get the new configuration.
- C. The checkdiskspace.timer should be configured to allow manual starts.
- D. The checkdiskspace.timer should be started using the sudo command.

Correct Answer: C

Explanation: The most likely reason the timer will not start is that the checkdiskspace.timer should be configured to allow manual starts. By default, systemd timers do not allow manual activation via systemctl start, unless they have `RefuseManualStart=no` in their [Unit] section. This option prevents users from accidentally starting timers that are meant to be controlled by other mechanisms, such as calendar events or dependencies. To enable manual starts for checkdiskspace.timer, the administrator should add `RefuseManualStart=no` to its [Unit] section and reload systemd. The other options are not correct reasons for the timer not starting. The checkdiskspace.timer unit does not need to be enabled via systemctl enable, because enabling a timer only makes it start automatically at boot time or after a system reload, but does not affect manual activation. The timers.target does not need to be reloaded to get the new configuration, because reloading a target only affects units that have a dependency on it, but does not affect manual activation. The checkdiskspace.timer does not need to be started using the sudo command, because the administrator is already running systemctl as root, as indicated by the # prompt. References: systemd.timer(5) - Linux manual page; systemctl(1) - Linux manual page

[Latest XK0-005 Dumps](#)

[XK0-005 PDF Dumps](#)

[XK0-005 Braindumps](#)