CompTIA

Exam Questions XK0-005

CompTIA Linux+ Certification Exam

NEW QUESTION 1
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

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.

NEW QUESTION 2
A user reported issues when trying to log in to a Linux server. The following outputs were received:
Given the outputs above. which of the following is the reason the user is una-ble to log in to the server?

A. User1 needs to set a long password.
B. User1 is in the incorrect group.
C. The user1 shell assignment incorrect.
D. The user1 password is expired.

Answer: D

Explanation:
The user1 password is expired. This can be inferred from the output of the chage -l user1 command, which shows the password expiration information for user1. The output shows that the password expired on 2020-10-01, and the account expired on 2020-10-08. This means that user1 cannot log in to the server unless the password and account are reactivated by the system administrator.
The other options are not correct based on the outputs above. User1 does not need to set a long password, because the output of the passwd -S user1 command shows that the password has a minimum length of 5 characters, which is met by user1's password. User1 is not in the incorrect group, because the output of the groups user1 command shows that user1 belongs to the app group, which is presumably the correct group for accessing the server. The user1 shell assignment is not incorrect, because the output of the grep user1
/etc/passwd command shows that user1 has /bin/bash as the default shell, which is a valid and common shell for Linux users.

NEW QUESTION 3
After starting an Apache web server, the administrator receives the following error:
Apr 23 localhost.localdomain httpd 4618] : (98) Address already in use: AH00072: make_sock: could not bind to address [: :]80
Which of the following commands should the administrator use to further trou-bleshoot this issue?

A. Ss
B. Ip
C. Dig
D. Nc

Answer: A

Explanation:
The ss command is used to display information about socket connections, such as the port number, state, and process ID. The error message indicates that the port 80 is already in use by another process, which prevents the Apache web server from binding to it. By using the ss command with the -l and -n options, the administrator can list all the listening sockets and their port numbers in numeric form, and identify which process is using the port 80. For example: ss -ln | grep :80.
The ip, dig, and nc commands are not relevant for this issue, as they are used for different purposes, such as configuring network interfaces, querying DNS records, and testing network connectivity.

NEW QUESTION 4
In which of the following filesystems are system logs commonly stored?

A. /var
B. /tmp
C. /etc
D. /opt

Answer: A

Explanation:
The filesystem that system logs are commonly stored in is /var. The /var filesystem is a directory that contains variable data files on Linux systems. Variable data files are files that are expected to grow in size over time, such as logs, caches, spools, and temporary files. The /var filesystem is separate from the / filesystem, which contains the essential system files, to prevent the / filesystem from being filled up by the variable data files. The system logs are files that record the events and activities of the system and its components, such as the kernel, the services, the applications, and the users. The system logs are useful for monitoring, troubleshooting, and auditing the system. The system logs are commonly stored in the /var/log directory, which is a subdirectory of the /var filesystem. The /var/log directory contains various log files, such as syslog, messages, dmesg, auth.log, and kern.log. The filesystem that system logs are commonly stored in is /var. This is the correct answer to the question. The other options are incorrect because they are not the filesystems that system logs are commonly stored in (/tmp, /etc, or /opt). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Managing Logging and Monitoring, page 487.

NEW QUESTION 5
A non-privileged user is attempting to use commands that require elevated account permissions, but the commands are not successful. Which of the following most likely needs to be updated?

A. /etc/passwd
B. /etc/shadow
C. /etc/sudoers
D. /etc/bashrc

Answer: C

Explanation:
The /etc/sudoers file is used to configure the sudo command, which allows non-privileged users to execute commands that require elevated account permissions1. The file contains a list of users and groups that are allowed to use sudo, and the commands they can run with it. The file also defines the security policy for sudo, such as whether a password is required, how long the sudo session lasts, and what environment variables are preserved or reset.
The /etc/passwd file is used to store information about the user accounts on the system, such as their username, user ID, home directory, and login shell. The /etc/shadow file is used to store the encrypted passwords for the user accounts, along with other information such as password expiration and aging. These files are not directly related to the sudo command, and updating them will not grant a user elevated account permissions.
The /etc/bashrc file is used to set up the environment for the bash shell, such as aliases, functions, variables, and options. This file is executed whenever a new bash shell is started, and it affects all users on the system. However, this file does not control the sudo command or its configuration, and updating it will not allow a user to use commands that require elevated account permissions.

NEW 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. ls -11 /etc
C. lshw —class disk
D. pvdisplay
E. rmdir /dev
F. dmesg

Answer: CF

Explanation:
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.

NEW QUESTION 7
During a security scan, the password of an SSH key file appeared to be too weak and was cracked. Which of the following commands would allow a user to choose a stronger password and set it on the existing SSH key file?

A. passwd
B. ssh
C. ssh-keygen
D. pwgen

Answer: C

Explanation:
The command that would allow a user to choose a stronger password and set it on the existing SSH key file is ssh-keygen -p -f <keyfile>. This command uses the ssh-keygen tool, which is used to generate, manage, and convert authentication keys for SSH. The -p option stands for passphrase, and it allows the user to change or remove the passphrase of an existing private key file. The -f option specifies the filename of the key file. The command will prompt the user for the old passphrase, and then for the new passphrase twice.
The other options are not correct commands for changing the password of an SSH key file. The passwd command is used to change the password of a user account on a Linux system, not an SSH key file. The ssh command is used to log in to a remote system using SSH, not to change the password of an SSH key file. The pwgen command is used to generate random passwords, not to change the password of an SSH key file.
References: ssh-keygen(1) - Linux manual page; How To: Change Passphrase for SSH Private Key - Unix Tutorial

NEW QUESTION 8
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

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 mysqld 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

NEW QUESTION 9
A cloud engineer is installing packages during VM provisioning. Which of the following should the engineer use to accomplish this task?

A. Cloud-init
B. Bash
C. Docker
D. Sidecar

Answer: A

Explanation:
The cloud engineer should use cloud-init to install packages during VM provisioning. Cloud-init is a tool that allows the customization of cloud instances at boot time. Cloud-init can perform various tasks, such as setting the hostname, creating users, installing packages, configuring network, and running scripts. Cloud-init can work with different cloud platforms and Linux distributions. This is the correct tool to accomplish the task. The other options are incorrect because they are either not suitable for cloud provisioning (Bash or Docker) or not a tool but a design pattern (Sidecar). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 563.

NEW QUESTION 10
To harden one of the servers, an administrator needs to remove the possibility of remote administrative login via the SSH service. Which of the following should the administrator do?

A. Add the line DenyUsers root to the /etc/hosts.deny file.
B. Set PermitRootLogin to no in the /etc/ssh/sshd_config file.
C. Add the line account required pam_nologi
D. so to the /etc/pam.d/sshd file.
E. Set PubKeyAuthentication to no in the /etc/ssh/ssh_config file.

Answer: B

Explanation:
The administrator should set PermitRootLogin to no in the /etc/ssh/sshd_config file to remove the possibility of remote administrative login via the SSH service. The PermitRootLogin directive controls whether the root user can log in using SSH. Setting it to no will deny any remote login attempts by the root user. This will harden the server and prevent unauthorized access. The administrator should also restart the sshd service after making the change. The other options are incorrect because they either do not affect the SSH service (/etc/hosts.deny or /etc/pam.d/sshd) or do not prevent remote administrative login (PubKeyAuthentication). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 413.

NEW QUESTION 10
An administrator needs to make an application change via a script that must be run only in console mode. Which of the following best represents the sequence the administrator should execute to accomplish this task?

A. systemct1 isolate multi-user.target sh script.shsystemct1 isolate graphical.target
B. systemct1 isolate graphical.target sh script.shsystemct1 isolate multi-user.target

C. sh script.shsystemct1 isolate multi-user.target systemct1 isolate graphical.target
D. systemct1 isolate multi-user.target systemct1 isolate graphical.targetsh script.sh

Answer: A

Explanation:
The correct answer is A. systemctl isolate multi-user.target sh script.sh systemctl isolate graphical.target
This sequence will allow the administrator to switch from the graphical mode to the console mode, run the script, and then switch back to the graphical mode.
The systemctl command is used to control the systemd system and service manager, which manages the boot targets and services on Linux systems. The isolate subcommand starts the unit specified on the command line and its dependencies and stops all others. The multi-user.target is a boot target that provides a text-based console login, while the graphical.target is a boot target that provides a graphical user interface. By using systemctl isolate, the administrator can change the boot target on the fly without rebooting the system.
The sh command is used to run a shell script, which is a file that contains a series of commands that can be executed by the shell. The script.sh is the name of the script that contains the application change that the administrator needs to make. By running sh script.sh, the administrator can execute the script in the console mode.
The other options are incorrect because:
* B. systemctl isolate graphical.target sh script.sh systemctl isolate multi-user.target
This sequence will switch from the console mode to the graphical mode, run the script, and then switch back to the console mode. This is not what the administrator wants to do, as the script must be run only in console mode.
* C. sh script.sh systemctl isolate multi-user.target systemctl isolate graphical.target
This sequence will run the script in the current mode, which may or may not be console mode, and then switch to console mode and back to graphical mode. This is not what the administrator wants to do, as the script must be run only in console mode.
* D. systemctl isolate multi-user.target systemctl isolate graphical.target sh script.sh
This sequence will switch from graphical mode to console mode and then back to graphical mode, without running the script at all. This is not what the administrator wants to do, as the script must be run only in console mode.
References:
? systemctl(1) - Linux manual page
? How to switch between the CLI and GUI on a Linux server
? How to PROPERLY boot into single user mode in RHEL/CentOS 7/8
? Changing Systemd Boot Target in Linux
? Exit Desktop to Terminal in Ubuntu 19.10

NEW QUESTION 13
A user generated a pair of private-public keys on a workstation. Which of the following commands will allow the user to upload the public key to a remote server and enable passwordless login?

A. scp ~/.ssh/id_rsa user@server:~/
B. rsync ~ /.ssh/ user@server:~/
C. ssh-add user server
D. ssh-copy-id user@server

Answer: D

Explanation:
The command ssh-copy-id user@server will allow the user to upload the public key to a remote server and enable passwordless login. The ssh-copy-id command is a tool for copying the public key to a remote server and appending it to the authorized_keys file, which is used for public key authentication. The command will also set the appropriate permissions on the remote server to ensure the security of the key. The command ssh-copy-id user@server will copy the public key of the user to the server and allow the user to log in without a password. This is the correct command to use for this task. The other options are incorrect because they either do not copy the public key (scp, rsync, or ssh-add) or do not use the correct syntax (scp ~/.ssh/id_rsa user@server:~/ instead of scp ~/.ssh/id_rsa.pub user@server:~/ or rsync ~ /.ssh/ user@server:~/ instead of rsync ~/.ssh/id_rsa.pub user@server:~/). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 410.

NEW QUESTION 16
Which of the following directories is the mount point in a UEFI system?

A. /sys/efi
B. /boot/efi
C. /efi
D. /etc/efi

Answer: B

Explanation:
The /boot/efi directory is the mount point in a UEFI system. This directory contains the EFI System Partition (ESP), which stores boot loaders and other files required by UEFI firmware. The /sys/efi directory does not exist by default in Linux systems. The /efi directory does not exist by default in Linux systems. The /etc/efi directory does not exist by default in Linux systems. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing the Linux Boot Process, page 398.

NEW QUESTION 19
A systems administrator is gathering information about a file type and the contents of a file. Which of the following commands should the administrator use to accomplish this task?

A. file filename
B. touch filename
C. grep filename
D. lsof filename

Answer: A

Explanation:

The file command is used to determine the type of a file by examining its contents. It can recognize many different formats, such as text, binary, executable, compressed, image, audio, video, etc. It can also display some additional information about the file, such as encoding, size, dimensions, etc12
References: 1: file(1) - Linux manual page 2: How to use the file command in Linux

NEW QUESTION 20
A Linux engineer needs to block an incoming connection from the IP address 2.2.2.2 to a secure shell server and ensure the originating IP address receives a response that a firewall is blocking the connection. Which of the following commands can be used to accomplish this task?

A. iptables -A INPUT -p tcp -- dport ssh -s 2.2.2.2 -j DROP
B. iptables -A INPUT -p tcp -- dport ssh -s 2.2.2.2 -j RETURN
C. iptables -A INPUT -p tcp -- dport ssh -s 2.2.2.2 -j REJECT
D. iptables -A INPUT -p tcp -- dport ssh -s 2.2.2.2 -j QUEUE

Answer: C

Explanation:
The REJECT target sends back an error packet to the source IP address, indicating that the connection is refused by the firewall. This is different from the DROP target, which silently discards the packet without any response. The RETURN target returns to the previous chain, which may or may not accept the connection. The QUEUE target passes the packet to a userspace application for further processing, which is not the desired outcome in this case.
References
? CompTIA Linux+ (XK0-005) Certification Study Guide, page 316
? iptables - ssh - access from specific ip only - Server Fault, answer by Eugene Ionichev

NEW QUESTION 21
A DevOps engineer wants to allow the same Kubernetes container configurations to be deployed in development, testing, and production environments. A key requirement is that the containers should be configured so that developers do not have to statically configure custom, environment-specific locations. Which of the following should the engineer use to meet this requirement?

A. Custom scheduler
B. Node affinity
C. Overlay network
D. Ambassador container

Answer: D

Explanation:
To allow the same Kubernetes container configurations to be deployed in different environments without statically configuring custom locations, the engineer can use an ambassador container (D). An ambassador container is a proxy container that handles communication between containers and external services. It can dynamically configure locations based on environment variables or other methods. The other options are not related to this requirement. References:
? [CompTIA Linux+ Study Guide], Chapter 11: Working with Containers, Section: Using Ambassador Containers
? [How to Use Ambassador Containers]

NEW QUESTION 24
While inspecting a recently compromised Linux system, the administrator identified a number of processes that should not have been running:

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
5545	joe	30	-10	5465	56465	8254	R	0.5	1.5	00:35.3	upload.sh
2567	joe	30	-10	6433	75544	9453	R	0.7	1.8	00:25.1	upload_passwd.sh
8634	joe	30	-10	3584	74537	6435	R	0.3	1.1	00:17.6	uploadpw.sh
4846	joe	30	-10	6426	63234	9683	R	0.8	1.9	00:22.2	upload_shadow.sh

Which of the following commands should the administrator use to terminate all of the identified processes?

A. pkill -9 -f "upload*.sh"
B. kill -9 "upload*.sh"
C. killall -9 -upload*.sh"
D. skill -9 "upload*.sh"

Answer: A

Explanation:
The pkill -9 -f “upload*.sh” command will terminate all of the identified processes. This command will send a SIGKILL signal (-9) to all processes whose full command line matches the pattern “upload*.sh” (-f). This signal will force the processes to terminate immediately without giving them a chance to clean up or save their state. The kill -9 “upload*.sh” command is invalid, as kill requires a process ID (PID), not a pattern. The killall -9 “upload*.sh” command is incorrect, as killall requires an exact process name, not a pattern. The skill -9 “upload*.sh” command is incorrect, as skill requires a username or a session ID (SID), not a pattern. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Memory and Process Execution, page 470.

NEW QUESTION 29
A developer has been unable to remove a particular data folder that a team no longer uses. The developer escalated the issue to the systems administrator. The following output was received:

```
# rmdir data/
rmdir: failed to remove 'data/': Operation not permitted
# rm -rf data/
rm: cannot remove 'data': Operation not permitted
# mv data/ mydata
mv: cannot move 'data/' to 'mydata': Operation not permitted
# cd data/
# cat > test.txt
bash: test.txt: Permission denied
```

Which of the following commands can be used to resolve this issue?

A. chgrp -R 755 data/
B. chmod -R 777 data/
C. chattr -R -i data/
D. chown -R data/

Answer: C

Explanation:
The command that can be used to resolve the issue of being unable to remove a particular data folder is chattr -R -i data/. This command will use the chattr utility to change file attributes on a Linux file system. The -R option means that chattr will recursively change attributes of directories and their contents. The -i option means that chattr will remove (unset) the immutable attribute from files or directories. When a file or directory has the immutable attribute set, it cannot be modified, deleted, or renamed.
The other options are not correct commands for resolving this issue. The chgrp -R 755 data/ command will change the group ownership of data/ and its contents recursively to 755, which is not a valid group name. The chgrp command is used to change group ownership of files or directories. The chmod -R 777 data/ command will change the file mode bits of data/ and its contents recursively to 777, which means that everyone can read, write, and execute them. However, this will not remove the immutable attribute, which prevents deletion or modification regardless of permissions. The chmod command is used to change file mode bits of files or directories. The chown -R data/ command is incomplete and will produce an error. The chown command is used to change the user and/or group ownership of files or directories, but it requires at least one argument besides the file name. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 7: Managing Disk Storage; chattr(1) - Linux manual page; chgrp(1) - Linux manual page; chmod(1) - Linux manual page; chown(1) - Linux manual page

NEW QUESTION 33
A junior administrator is trying to set up a passwordless SSH connection to one of the servers. The administrator follows the instructions and puts the key in the authorized_key file at the server, but the administrator is still asked to provide a password during the connection.
Given the following output:

```
junior@server:-$ ls -lh .ssh/auth*
-rw------- 1 junior junior 566 sep 13 20:56 .ssh/authorized_key
```

Which of the following commands would resolve the issue and allow an SSH connection to be established without a password?

A. restorecon -rv .ssh/authorized_key
B. mv .ssh/authorized_key .ssh/authorized_keys
C. systemct1 restart sshd.service
D. chmod 600 mv .ssh/authorized_key

Answer: B

Explanation:
The command mv .ssh/authorized_key .ssh/authorized_keys will resolve the issue and allow an SSH connection to be established without a password. The issue is caused by the incorrect file name of the authorized key file on the server. The file should be named authorized_keys, not authorized_key. The mv command will rename the file and fix the issue. The other options are incorrect because they either do not affect the file name (restorecon or chmod) or do not restart the SSH service (systemct1). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 410.

NEW QUESTION 38
A Linux systems administrator receives reports from various users that an application hosted on a server has stopped responding at similar times for several days in a row. The administrator logs in to the system and obtains the following output:
Output 1:

```
[Tue Aug 31 16:36:42 2021] OOM: Kill process 43805 (java) score 249 or sacrifice child
[Tue Aug 31 16:36:42 2021] killed process 43805 (java) total-vm: 4446352kB, anon-rss: 4053140kB, file-rss: 68kB
```

Output 2:

```
Linux 3.10.0-328.13.1.x86_64 #1 (hostname) 31/08/2021 _x86_64_ (8 CPU)
16:00:01 PM     CPU    %user   %nice   %system   %iowait   %steal       %idle
16:10:01 PM     all    17.58    0.00      9.36      0.00     0.00       73.06
16:20:01 PM     all    22.34    0.00     11.75      0.00     0.00       65.91
16:30:01 PM     all    25.49    0.00     11.69      0.00        0       62.82
```

Output 3:

```
$ free -m
              total        used        free      shared  buff/cache   available
Mem:          16704       15026         174          92         619         793
Swap:             0           0           0
```

Which of the following should the administrator do to provide the BEST solution for the reported issue?

A. Configure memory allocation policies during business hours and prevent the Java process from going into a zombie state while the server is idle.
B. Configure a different nice value for the Java process to allow for more users and prevent the Java process from restarting during business hours.
C. Configure more CPU cores to allow for the server to allocate more processing and prevent the Java process from consuming all of the available resources.
D. Configure the swap space to allow for spikes in usage during peak hours and prevent the Java process from stopping due to a lack of memory.

Answer: D

Explanation:
Based on the output of the image sent by the user, the system requires more swap space to allow for spikes in usage during peak hours and prevent the Java process from stopping due to a lack of memory. The output shows that there is only 0 MB of swap space available on the system, which means that there is no room for swapping out memory pages when physical memory is full or low. The output also shows that there is only 793 MB of available memory on the system, which may not be enough to handle high- demand applications such as Java. This may cause Java to stop working due to insufficient memory or trigger an OutOfMemoryError exception. Configuring more swap space on the system would help to alleviate this issue by providing more virtual memory for applications and improving performance. Configuring memory allocation policies during business hours will not help to solve this issue, as it will not increase the amount of available memory or swap space on the system. Configuring a different nice value for Java process will not help to solve this issue, as it will only affect its scheduling priority, not its memory consumption or allocation. Configuring more CPU cores will not help to solve this issue, as it will only increase processing power, not memory capacity or availability. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Memory and Process Execution, page 468.

NEW QUESTION 41
Which of the following is the best tool for dynamic tuning of kernel parameters?

A. tuned
B. tune2fs
C. tuned-adm
D. turbostat

Answer: A

Explanation:
The tuned application is the best tool for dynamic tuning of kernel parameters, as it monitors the system and optimizes the performance under different workloads. It provides a number of predefined profiles for typical use cases, such as power saving, low latency, high throughput, virtual machine performance, and so on. It also allows users to create, modify, and delete profiles, and to switch between them on the fly. The tuned application uses the sysctl command and the configuration files in the /etc/sysctl.d/ directory to adjust the kernel parameters at runtime.
References
? Chapter 2. Getting started with TuneD - Red Hat Customer Portal, paragraph 1
? Kernel tuning with sysctl - Linux.com, paragraph 1

NEW QUESTION 44
A junior administrator is setting up a new Linux server that is intended to be used as a router at a remote site. Which of the following parameters will accomplish this goal?
A.

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A PREROUTING -i eth0 -j MASQUERADE
```

A.

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -D POSTROUTING -o eth0 -j MASQUERADE
```

B.

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
```

C.

```
echo 1 > /proc/sys/net/ipv4/ip_forward
iptables -t nat -A PREROUTING -o eth0 -j MASQUERADE
```

Answer: C

Explanation:
The parameter net.ipv4.ip_forward=1 will accomplish the goal of setting up a new Linux server as a router. This parameter enables the IP forwarding feature, which allows the server to forward packets between different network interfaces. This is necessary for a router to route traffic between different networks. The parameter can be set
in the /etc/sysctl.conf file or by using the sysctl command. This is the correct parameter to use to accomplish the goal. The other options are incorrect because they either do not exist (net.ipv4.ip_forwarding or net.ipv4.ip_route) or do not enable IP forwarding (net.ipv4.ip_forward=0). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 382.

NEW QUESTION 49
A systems administrator is tasked with preventing logins from accounts other than root, while the file /etc/nologin exists. Which of the following PAM modules will accomplish this task?

A. pam_login.so
B. pam_access.so
C. pam_logindef.so
D. pam_nologin.so

Answer: D

Explanation:
The PAM module pam_nologin.so will prevent logins from accounts other than root, while the file /etc/nologin exists. This module checks for the existence of the file /etc/nologin and displays its contents to the user before denying access. The root user is exempt from this check and can still log in. This is the correct module to accomplish the task. The other options are incorrect because they are either non-existent modules (pam_login.so or pam_logindef.so) or do not perform the required function (pam_access.so controls access based on host, user, or time). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Users and Groups, page 471.

NEW QUESTION 54
A systems administrator is receiving tickets from users who cannot reach the application app that should be listening on port 9443/tcp on a Linux server.
To troubleshoot the issue, the systems administrator runs netstat and receives the following output:

```
# netstat -anp | grep appd | grep -w LISTEN
tcp 0 0 127.0.0.1:9443 0.0.0.0:* LISTEN 1234/appd
```

Based on the information above, which of the following is causing the issue?

A. The IP address 0.0.0.0 is not valid.
B. The application is listening on the loopback interface.
C. The application is listening on port 1234.
D. The application is not running.

Answer: B

Explanation:
The server is in a "Listen" state on port 9943 using its loopback address. The "1234" is a process-id
The cause of the issue is that the application is listening on the loopback interface. The loopback interface is a virtual network interface that is used for internal communication within the system. The loopback interface has the IP address 127.0.0.1, which is also known as localhost. The netstat output shows that the application is listening on port 9443 using the IP address 127.0.0.1. This means that the application can only accept connections from the same system, not from other systems on the network. This can prevent the users from reaching the application and cause the issue. The administrator should configure the application to listen on the IP address 0.0.0.0, which means all available interfaces, or on the specific IP address of the system that is reachable from the network. This will allow the application to accept connections from other systems and resolve the issue. The cause of the issue is that the application is listening on the loopback interface. This is the correct answer to the question. The other options are incorrect because they are not supported by the outputs. The IP address 0.0.0.0 is valid and means all interfaces, the application is not listening on port 1234, and the application is running as shown by the process ID 1234. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 383.

NEW QUESTION 56
A development team asks an engineer to guarantee the persistency of journal log files across system reboots. Which of the following commands would accomplish this task?

A. grep -i auto /etc/systemd/journald.conf && systemct1 restart systemd-journald.service
B. cat /etc/systemd/journald.conf | awk '(print $1,$3)'
C. sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i 'persistent/s/^#//q'/etc/systemd/journald.conf
D. journalctl --list-boots && systemct1 restart systemd-journald.service

Answer: C

Explanation:
The command sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed - i 'persistent/s/^#//q' /etc/systemd/journald.conf will accomplish the task of guaranteeing the persistency of journal log files across system reboots. The sed command is a tool for editing text files on Linux systems. The -i option modifies the file in place. The s command substitutes one string for another. The g flag replaces all occurrences of the string.
The && operator executes the second command only if the first command succeeds. The q command quits after the first match. The /etc/systemd/journald.conf file is a configuration file for the systemd-journald service, which is responsible for collecting and storing log messages. The command sed -i 's/auto/persistent/g' /etc/systemd/journald.conf will replace the word auto with the word persistent in the file. This will change the value of the Storage option, which controls where the journal log files are stored. The value auto means that the journal log files are stored in the volatile memory and are lost after reboot, while the value persistent means that the journal log files are stored in the persistent storage and are preserved across reboots. The command sed -i 'persistent/s/^#//q' /etc/systemd/journald.conf will remove the # character at the beginning of the line that contains the word persistent. This will uncomment the Storage option and enable it. The command sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i 'persistent/s/^#//q' /etc/systemd/journald.conf will guarantee the persistency of journal log files across system reboots by changing and enabling the Storage option to persistent. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not change the value of the Storage option (grep -i auto /etc/systemd/journald.conf && systemct1 restart systemd-journald.service or cat /etc/systemd/journald.conf | awk '(print $1,$3)') or do not enable the Storage option (journalctl --list-boots && systemct1 restart systemd- journald.service). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Managing Logging and Monitoring, page 489.

NEW QUESTION 58
A Linux administrator is adding a new configuration file to a Git repository. Which of the following describes the correct order of Git commands to accomplish the task successfully?

A. pull -> push -> add -> checkout
B. pull -> add -> commit -> push
C. checkout -> push -> add -> pull
D. pull -> add -> push -> commit

Answer: B

Explanation:
The correct order of Git commands to add a new configuration file to a Git repository is pull -> add -> commit -> push. The pull command will fetch and merge the changes from the remote repository to the local repository, ensuring that the local repository is up to date. The add command will stage the new configuration file for the next commit, marking it as a new file to be tracked by Git. The commit command will create a new snapshot of the project state with the new configuration file and a descriptive message. The push command will publish the commit to the remote repository, updating the remote branch with the new configuration file. The pull -> push -> add -> checkout order is incorrect, as it will not create a commit for the new configuration file, and it will switch to a different branch without pushing the changes. The checkout -> push -> add -> pull order is incorrect, as it will switch to a different branch before adding the new configuration file, and it will overwrite the local changes with the remote changes without creating a commit. The pull -> add -> push -> commit order is incorrect, as it will not create a commit before pushing the changes, and it will create a commit that is not synchronized with the remote branch. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 20: Writing and Executing Bash Shell Scripts, page 579.

NEW QUESTION 60
DRAG DROP
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 /.

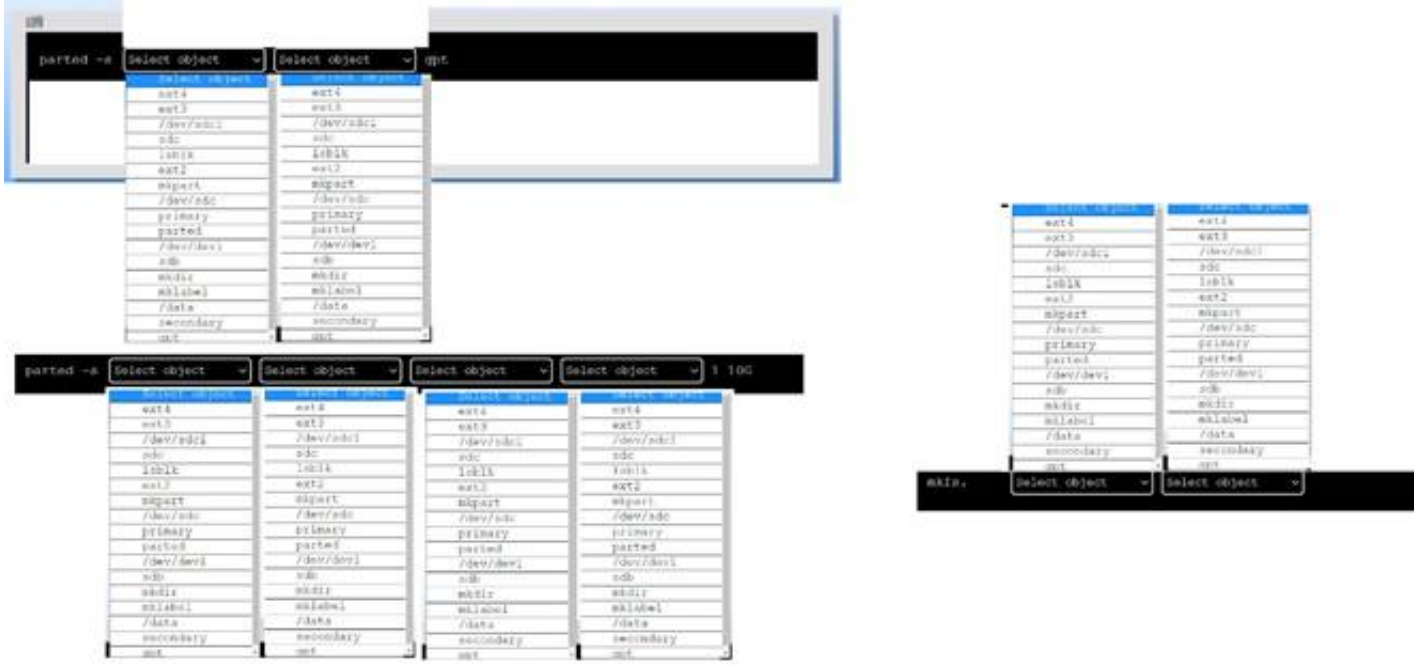

A. Mastered
B. Not Mastered

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 61
A Linux engineer has been notified about the possible deletion of logs from the file
/opt/app/logs. The engineer needs to ensure the log file can only be written into without removing previous entries.

```
# lsattz /opt/app/logs
--------------e---  logs
```

Which of the following commands would be BEST to use to accomplish this task?

A. chattr +a /opt/app/logs
B. chattr +d /opt/app/logs
C. chattr +i /opt/app/logs

D. chattr +c /opt/app/logs

Answer: A

Explanation:
The command chattr +a /opt/app/logs will ensure the log file can only be written into without removing previous entries. The chattr command is a tool for changing file attributes on Linux file systems. The +a option sets the append-only attribute, which means that the file can only be opened in append mode for writing. This prevents the file from being modified, deleted, or renamed. This is the best command to use to accomplish the task. The other options are incorrect because they either set the wrong attributes
(+d, +i, or +c) or do not affect the file at all (-a). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, page 357.

NEW QUESTION 66
Users are unable to create new files on the company's FTP server, and an administrator is troubleshooting the issue. The administrator runs the following commands:

```
# df -h /ftpusers/

Filesystem      Size      Used      Avail     Use%     Mounted on
/dev/sda4       150G      40G       109G      26%      /ftpusers

# df -i /ftpusers/

Filesystem      Inodes    Iused     Ifree     Iuse%    Mounted on
/dev/sda4       34567     34567     0         100%     /ftpusers
```

Which of the following is the cause of the issue based on the output above?

A. The users do not have the correct permissions to create files on the FTP server.
B. The ftpusers filesystem does not have enough space.
C. The inodes is at full capacity and would affect file creation for users.
D. ftpusers is mounted as read only.

Answer: C

Explanation:
The cause of the issue based on the output above is C. The inodes is at full capacity and would affect file creation for users.
An inode is a data structure that stores information about a file or directory, such as its name, size, permissions, owner, timestamps, and location on the disk. Each file or directory has a unique inode number that identifies it. The number of inodes on a filesystem is fixed when the filesystem is created, and it determines how many files and directories can be created on that filesystem. If the inodes are exhausted, no new files or directories can be created, even if there is enough disk space available.
The output for the second command shows that the /ftpusers/ filesystem has 0% of inodes available, which means that all the inodes have been used up. This would prevent users from creating new files on the FTP server. The administrator should either delete some unused files or directories to free up some inodes, or resize the filesystem to increase the number of inodes.
The other options are incorrect because:
* A. The users do not have the correct permissions to create files on the FTP server.
This is not true, because the output for the first command shows that the /ftpusers/ filesystem has 26% of disk space available, which means that there is enough space for users to create files. The permissions of the files and directories are not shown in the output, but they are not relevant to the issue of inode exhaustion.
* B. The ftpusers filesystem does not have enough space.
This is not true, because the output for the first command shows that the /ftpusers/ filesystem has 26% of disk space available, which means that there is enough space for users to create files. The issue is not related to disk space, but to inode capacity.
* D. ftpusers is mounted as read only.
This is not true, because the output for the first command does not show any indication that the /ftpusers/ filesystem is mounted as read only. If it was, it would have an (ro) flag next to the mounted on column. A read only filesystem would prevent users from creating or modifying files on the FTP server, but it would not affect the inode usage.

NEW QUESTION 71
A systems administrator is encountering performance issues. The administrator runs 3 commands with the following output

```
09:10:18  up  457 days,  32min,  5 users,  load average:  4.22  6.63  5.98
```

The Linux server has the following system properties CPU: 4 vCPU
Memory: 50GB
Which of the following accurately describes this situation?

A. The system is under CPU pressure and will require additional vCPUs
B. The system has been running for over a year and requires a reboot.
C. Too many users are currently logged in to the system
D. The system requires more memory

Answer: A

Explanation:
Based on the output of the image sent by the user, the system is under CPU pressure and will require additional vCPUs. The output shows that there are four processes running upload.sh scripts that are consuming a high percentage of CPU time (99.7%, 99.6%, 99.5%, and 99.4%). The output also shows that the system has only 4 vCPUs, which means that each process is using almost one entire vCPU. This indicates that the system is struggling to handle the CPU load and may experience performance issues or slowdowns. Adding more vCPUs to the system would help to alleviate the CPU pressure and improve the system performance. The system has not been running for over a year, as the uptime command shows that it has been up for only 1 day, 2 hours, and 13 minutes. The

number of users logged in to the system is not relevant to the performance issue, as they are not consuming significant CPU resources. The system does not require more memory, as the free command shows that it has plenty of available memory (49 GB total, 48 GB free). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Memory and Process Execution, pages 468-469.

NEW QUESTION 76
A systems administrator wants to check for running containers. Which of the following commands can be used to show this information?

A. docker pull
B. docker stats
C. docker ps
D. docker list

Answer: C

Explanation:
The command that can be used to check for running containers is docker ps. The docker ps command can list all the containers that are currently running on the system. To show all the containers, including those that are stopped, the administrator can use docker ps -a
R©e. ferences:
? [CompTIA Linux+ Study Guide], Chapter 11: Working with Containers, Section: Managing Containers with Docker
? [Docker PS Command with Examples]

NEW QUESTION 77
A systems administrator was tasked with assigning the temporary IP address/netmask 192.168.168.1/255.255.255.255 to the interface eth0 of a Linux server. When adding the address, the following error appears:
ip address add 192.168.168.1/33 dev eth0
Error: any valid prefix is expected rather than "192.168.168.1/33".
Based on the command and its output above, which of the following is the cause of the issue?

A. The CIDR value /33 should be /32 instead.
B. There is no route to 192.168.168.1/33.
C. The interface eth0 does not exist.
D. The IP address 192.168.168.1 is already in use.

Answer: A

Explanation:
The cause of the issue is that the CIDR value /33 is invalid for an IPv4 address. The CIDR value represents the number of bits in the network prefix of an IP address, and it can range from 0 to 32 for IPv4 addresses. A CIDR value of /33 would imply a network prefix of more than 32 bits, which is impossible for an IPv4 address. To assign a temporary IP address/netmask of 192.168.168.1/255.255.255.255 to eth0, the CIDR value should be /32 instead, which means a network prefix of 32 bits and a host prefix of 0 bits. There is no route to 192.168.168.1/33 is not the cause of the issue, as the ip address add command does not check the routing table. The interface eth0 does not exist is not the cause of the issue, as the ip address add command would display a different error message if the interface does not exist. The IP address 192.168.168.1 is already in use is not the cause of the issue, as the ip address add command would display a different error message if the IP address is already in use. References: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 13: Networking Fundamentals, page 435.

NEW QUESTION 81
A Linux administrator is creating a primary partition on the replacement hard drive for an application server. Which of the following commands should the administrator issue to verify the device name of this partition?

A. sudo fdisk /dev/sda
B. sudo fdisk -s /dev/sda
C. sudo fdisk -l
D. sudo fdisk -h

Answer: C

Explanation:
The command sudo fdisk -l should be issued to verify the device name of the partition. The sudo command allows the administrator to run commands as the superuser or another user. The fdisk command is a tool for manipulating disk partitions on Linux systems. The -l option lists the partitions on all disks or a specific disk. The command sudo fdisk -l will show the device names, sizes, types, and other information of the partitions on all disks. The administrator can identify the device name of the partition by looking at the output. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not list the partitions (sudo fdisk /dev/sda or sudo fdisk -h) or do not exist (sudo fdisk -s /dev/sda). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 317.

NEW QUESTION 83
One leg of an LVM-mirrored volume failed due to the underlying physical volume, and a systems administrator is troubleshooting the issue. The following output has been provided:

```
Partial mode. Incomplete volume groups will be activated read-only
```

LV	VG	Attr	LSize	Origin	Snap%	Move	Log	Copy%	Devices
linear	vg	-wi-a-	40.00G						unknown device(0)
stripe	vg	-wi-a-	40.00G						unknown device(5120),/dev/sda1(0)

Given this scenario, which of the following should the administrator do to recover this volume?

A. Reboot the serve
B. The volume will automatically go back to linear mode.
C. Replace the failed drive and reconfigure the mirror.
D. Reboot the serve
E. The volume will revert to stripe mode.

F. Recreate the logical volume.

Answer: B

Explanation:
The administrator should replace the failed drive and reconfigure the mirror to recover the volume. The LVM (Logical Volume Manager) is a tool for managing disk space on Linux systems. The LVM allows the administrator to create logical volumes that span across multiple physical volumes, such as hard disks or partitions. The LVM also supports different types of logical volumes, such as linear, striped, or mirrored. A mirrored logical volume is a type of logical volume that creates a copy of the data on another physical volume, providing redundancy and fault tolerance. The output shows that the logical volume is mirrored and that one leg of the mirror has failed due to the underlying physical volume. This means that one of the physical volumes that contains the data of the logical volume is damaged or missing. This can cause data loss and performance degradation. The administrator should replace the failed drive and reconfigure the mirror to recover the volume. The administrator should identify the failed physical volume by using commands such as pvdisplay, vgdisplay, or lvdisplay. The administrator should then remove the failed physical volume from the volume group by using the vgreduce command.
The administrator should then install a new drive and create a new physical volume by using the pvcreate command. The administrator should then add the new physical volume to the volume group by using the vgextend command. The administrator should then reconfigure the mirror by using the lvconvert command. The administrator should replace the failed drive and reconfigure the mirror to recover the volume. This is the correct answer to the question. The other options are incorrect because they either do not recover the volume (reboot the server. The volume will automatically go back to linear mode or reboot the server. The volume will revert to stripe mode) or do not preserve the data of the volume (recreate the logical volume). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 333-334.

NEW QUESTION 87
A Linux administrator recently downloaded a software package that is currently in a compressed file. Which of the following commands will extract the files?

A. unzip -v
B. bzip2 -z
C. gzip
D. funzip

Answer: C

Explanation:
The command gzip can extract files that are compressed with the gzip format, which has the extension .gz. This is the correct command to use for the software package. The other options are incorrect because they either compress files (bzip2 -z), unzip files that are compressed with the zip format (unzip -v or funzip), or have the wrong options (-v or -z instead of -d). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, page 353.

NEW QUESTION 91
Users have reported that the interactive sessions were lost on a Linux server. A Linux administrator verifies the server was switched to rescue.target mode for maintenance. Which of the following commands will restore the server to its usual target?

A. telinit 0
B. systemct1 reboot
C. systemct1 get-default
D. systemct1 emergency

Answer: B

Explanation:
The systemct1 reboot command will restore the server to its usual target by rebooting it. This will cause the server to load the default target specified in /etc/systemd/system.conf or /etc/systemd/system/default.target files. The telinit 0 command would shut down the server, not restore it to its usual target. The systemct1 get-default command would display the default target, not change it. The systemct1 emergency command would switch the server to emergency.target mode, which is even more
restrictive than rescue.target mode. References: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 17: System Maintenance and Operation, page 516.

NEW QUESTION 95
A systems administrator is tasked with changing the default shell of a system account in order to disable iterative logins. Which of the following is the best option for the administrator to use as the new shell?

A. /sbin/nologin
B. /bin/ sh
C. /sbin/ setenforce
D. /bin/bash

Answer: A

Explanation:
The /sbin/nologin shell is a special shell that prevents the user from logging into an interactive session. It is commonly used for system accounts that are not meant to be accessed by users, such as daemon or service accounts. When a user tries to log in with this shell, they will see a message like “This account is currently not available” and the login will fail.
References:
? The /sbin/nologin shell is listed as one of the valid shells in the /etc/shells file1.
? The CompTIA Linux+ Certification Exam Objectives mention that the candidate should be able to “configure and manage system accounts and groups, including password aging and restricted shells” as part of the Hardware and System Configuration domain2.
? The usermod command can be used to change the user’s login shell with the -s or --shell option3. For example, to change the shell of a user named daemon to /sbin/nologin, the command would be: sudo usermod -s /sbin/nologin daemon

NEW QUESTION 97
A systems administrator is installing various software packages using a pack-age manager. Which of the following commands would the administrator use on the

Linux server to install the package?

A. winget
B. softwareupdate
C. yum-config
D. apt

Answer: D

NEW QUESTION 100
A systems administrator is implementing a new service task with systems at startup and needs to execute a script entitled test.sh with the following content:

```
TIMESTAMP=$ (date '+%Y-%m-%d %H:%M:%S')
echo "helpme.service: timestamp $(Timestamp)" | systemd-cat -p info
sleep 60
done
```

The administrator tries to run the script after making it executable with chmod +x; however, the script will not run. Which of the following should the administrator do to address this issue? (Choose two.)

A. Add #!/bin/bash to the bottom of the script.
B. Create a unit file for the new service in /etc/systemd/system/ with the name helpme.service in the location.
C. Add #!//bin/bash to the top of the script.
D. Restart the computer to enable the new service.
E. Create a unit file for the new service in /etc/init.d with the name helpme.service in the location.
F. Shut down the computer to enable the new service.

Answer: BC

Explanation:
The administrator should do the following two things to address the issue:
? Add #!/bin/bash to the top of the script. This is called a shebang line and it tells the system which interpreter to use to execute the script. Without this line, the script will not run properly. The shebang line should be the first line of the script and should start with #! followed by the path to the interpreter. In this case, the interpreter is bash and the path is /bin/bash. The other option (A) is incorrect because the shebang line should be at the top, not the bottom of the script.
? Create a unit file for the new service in /etc/systemd/system/ with the name helpme.service in the location. This is necessary to register the script as a systemd service and enable it to run at startup. A unit file is a configuration file that defines the properties and behavior of a service, such as the description, dependencies, start and stop commands, and environment variables. The unit file should have the extension .service and should be placed in the /etc/systemd/system/ directory. The other option (E) is incorrect because /etc/init.d is the directory for init scripts, not systemd services.
References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, pages 429-430.

NEW QUESTION 103
A DevOps engineer needs to allow incoming traffic to ports in the range of 4000 to 5000 on a Linux server. Which of the following commands will enforce this rule?

A. iptables -f filter -I INPUT -p tcp --dport 4000:5000 -A ACCEPT
B. iptables -t filter -A INPUT -p tcp --dport 4000:5000 -j ACCEPT
C. iptables filter -A INPUT -p tcp --dport 4000:5000 -D ACCEPT
D. iptables filter -S INPUT -p tcp --dport 4000:5000 -A ACCEPT

Answer: B

Explanation:
The command iptables -t filter -A INPUT -p tcp --dport 4000:5000 -j ACCEPT will enforce the rule of allowing incoming traffic to ports in the range of 4000 to 5000 on a Linux server. The iptables command is a tool for managing firewall rules on Linux systems. The -t option specifies the table to operate on, in this case filter, which is the default table that contains the rules for filtering packets. The -A option appends a new rule to the end of a chain, in this case INPUT, which is the chain that processes the packets that are destined for the local system. The -p option specifies the protocol to match, in this case tcp, which is the transmission control protocol. The --dport option specifies the destination port or port range to match, in this case 4000:5000, which is the range of ports from 4000 to 5000. The -j option specifies the target to jump to if the rule matches, in this case ACCEPT, which is the target that allows the packet to pass through.
The command iptables -t filter -A INPUT -p tcp --dport 4000:5000 -j ACCEPT will add a new rule to the end of the INPUT chain that will accept the incoming TCP packets that have a destination port between 4000 and 5000. This command will enforce the rule and allow the traffic to the specified ports. This is the correct command to use to accomplish the task. The other options are incorrect because they either use the wrong options (-f instead of - t or -D instead of -A) or do not exist (iptables filter -A INPUT -p tcp --dport 4000:5000 -D ACCEPT or iptables filter -S INPUT -p tcp --dport 4000:5000 -A ACCEPT). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 543.

NEW QUESTION 108
A database administrator requested the installation of a custom database on one of the servers. Which of the following should the Linux administrator configure so the requested packages can be installed?

A. /etc/yum.conf
B. /etc/ssh/sshd.conf
C. /etc/yum.repos.d/db.repo
D. /etc/resolv.conf

Answer: C

Explanation:
The Linux administrator should configure /etc/yum.repos.d/db.repo so that the requested packages can be installed. This file defines a custom repository for yum, which is a package manager for RPM-based systems. The file should contain information such as the name, baseurl, gpgcheck, and enabled options for the repository. By creating this file and enabling the repository, the administrator can use yum to install packages from the custom repository. The /etc/yum.conf file is the main configuration file for yum, but it does not define repositories. The /etc/ssh/sshd.conf file is the configuration file for sshd, which is a daemon that provides

secure shell access to remote systems. The /etc/resolv.conf file is the configuration file for DNS resolution, which maps domain names to IP addresses. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Packages and Software, page 559.

NEW QUESTION 109
A Linux administrator is configuring a two-node cluster and needs to be able to connect the nodes to each other using SSH keys from the root account. Which of the following commands will accomplish this task?

A. [root@nodea ssh —i ~/ . ssh/±d rsa root@nodeb
B. [root@nodea scp -i . ssh/id rsa root@nodeb
C. [root@nodea ssh—copy-id —i .ssh/id rsa root@nodeb
D. [root@nodea # ssh add -c ~/ . ssh/id rsa root@nodeb
E. [root@nodea # ssh add -c ~/. ssh/id rsa root@nodeb

Answer: C

Explanation:
The ssh-copy-id command is used to copy a public SSH key from a local machine to a remote server and add it to the authorized_keys file, which allows passwordless authentication between the machines. The administrator can use this command to copy the root user's public key from nodea to nodeb, and vice versa, to enable SSH access between the nodes without entering a password every time. For example: [root@nodea ssh-copy-id -i ~/.ssh/id_rsa root@nodeb]. The ssh command is used to initiate an SSH connection to a remote server, but it does not copy any keys. The scp command is used to copy files securely between machines using SSH, but it does not add any keys to the authorized_keys file. The ssh-add command is used to add private keys to the SSH agent, which manages them for SSH authentication, but it does not copy any keys to a remote server.

NEW QUESTION 114
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

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.

NEW QUESTION 117
A Linux administrator is creating a new sudo profile for the accounting user. Which of the following should be added by the administrator to the sudo configuration file so that the accounting user can run / opt/ acc/ report as root?

A. accounting localhost=/opt/acc/report
B. accounting ALL=/opt/acc/report
C. %accounting ALL=(ALL) NOPASSWD: /opt/acc/report
D. accounting /opt/acc/report= (ALL) NOPASSWD: ALL

Answer: C

Explanation:
This answer allows the accounting user to run the /opt/acc/report command as root on any host without entering a password. The % sign indicates that accounting is a group name, not a user name. The ALL keyword means any host, any user, and any command, depending on the context. The NOPASSWD tag overrides the default behavior of sudo, which is to ask for the user's password.
The other answers are incorrect for the following reasons:
? A. accounting localhost=/opt/acc/report
? B. accounting ALL=/opt/acc/report
? D. accounting /opt/acc/report= (ALL) NOPASSWD: ALL

NEW QUESTION 122
The applications team is reporting issues when trying to access the web service hosted in a Linux system. The Linux systems administrator is reviewing the following outputs:
Output 1:
* httpd.service = The Apache HTTPD Server
Loaded: loaded (/usr/lib/systemd/system/httpd.service; disabled; vendor preset: disabled) Active: inactive (dead)
Docs: man:httpd(8) man:apachectl(8) Output 2:
16:51:16 up 28 min, 1 user, load average: 0.00, 0.00, 0.07
Which of the following statements best describe the root cause? (Select two).

A. The httpd service is currently started.
B. The httpd service is enabled to auto start at boot time, but it failed to start.
C. The httpd service was manually stopped.
D. The httpd service is not enabled to auto start at boot time.
E. The httpd service runs without problems.
F. The httpd service did not start during the last server reboot.

Answer: CD

Explanation:

The httpd.service is the Apache HTTPD Server, which is a web service that runs on Linux systems. The output 1 shows that the httpd.service is inactive (dead), which means that it is not running. The output 1 also shows that the httpd.service is disabled, which means that it is not enabled to auto start at boot time. Therefore, the statements C and D best describe the root cause of the issue. The statements A, B, E, and F are incorrect because they do not match the output 1. References: [How to Manage Systemd Services on a Linux System]

NEW QUESTION 125
A systems administrator has been tasked with disabling the nginx service from the environment to prevent it from being automatically and manually started. Which of the following commands will accomplish this task?

A. systemct1 cancel nginx
B. systemct1 disable nginx
C. systemct1 mask nginx
D. systemct1 stop nginx

Answer: C

Explanation:
The command systemct1 mask nginx disables the nginx service from the environment and prevents it from being automatically and manually started. This command creates a symbolic link from the service unit file to /dev/null, which makes the service impossible to start. This is the correct way to accomplish the task. The other options are incorrect because they either do not exist (systemct1 cancel nginx), do not prevent manual start (systemct1 disable nginx), or do not prevent automatic start (systemct1 stop nginx). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, page 429.

NEW QUESTION 130
A systems administrator is configuring a Linux system so the network traffic from the internal network 172.17.0.0/16 going out through the eth0 interface would appear as if it was sent directly from this interface. Which of the following commands will accomplish this task?

A. iptables -A POSTROUTING -s 172.17.0.0/16 -o eth0 -j MASQUERADE
B. firewalld -A OUTPUT -s 172.17.0.0/16 -o eth0 -j DIRECT
C. nmcli masq-traffic eth0 -s 172.17.0.0/16 -j MASQUERADE
D. ifconfig -- nat eth0 -s 172.17.0.0/16 -j DIRECT

Answer: A

Explanation:
This command will use the iptables tool to append a rule to the POSTROUTING chain of the nat table, which will match any packet with a source address of 172.17.0.0/16 and an output interface of eth0, and apply the MASQUERADE target to it. This means that the packet will have its source address changed to the address of the eth0 interface, effectively hiding the internal network behind a NAT12.
References: 1: Iptables NAT and Masquerade rules - what do they do? 2: Routing from docker containers using a different physical network interface and default gateway

NEW QUESTION 133
A Linux administrator modified the SSH configuration file. Which of the following commands should be used to apply the configuration changes?

A. systemct1 stop sshd
B. systemct1 mask sshd
C. systemct1 reload sshd
D. systemct1 start sshd

Answer: C

Explanation:
The systemct1 reload sshd command can be used to apply the configuration changes of the SSH server daemon without restarting it. This is useful to avoid interrupting existing connections. The systemct1 stop sshd command would stop the SSH server daemon, not apply the changes. The systemct1 mask sshd command would prevent the SSH server daemon from being started, not apply the changes. The systemct1 start sshd command would start the SSH server daemon if it is not running, but it would not apply the changes if it is already running. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Secure Shell (SSH), page 415.

NEW QUESTION 135
An application developer received a file with the following content:
##This is a sample Image ## FROM ubuntu:18.04
MAINTAINER demohut@gtmail.com.hac COPY . /app
RUN make /app
CMD python /app/app.py RUN apt-get update
RUN apt-get install -y nginx CMD ["echo","Image created"]
The developer must use this information to create a test bed environment and identify the image (myimage) as the first version for testing a new application before moving it to production. Which of the following commands will accomplish this task?

A. docker build -t myimage:1.0 .
B. docker build -t myimage: .
C. docker build -t myimage-1.0 .
D. docker build -i myimage:1.0 .

Answer: A

Explanation:
The docker build command is used to build an image from a Dockerfile and a context1. The Dockerfile is a text file that contains the instructions for creating the image, and the context is a set of files that can be used in the image creation process1. The file that the developer received is an example of a Dockerfile.
The -t option is used to specify a name and an optional tag for the image1. The name and tag are separated by a colon (:), and the tag is usually used to indicate

the version of the image2. For example, -t myimage:1.0 means that the image will be named myimage and tagged as 1.0.
The last argument of the docker build command is the path to the context, which can be a local directory or a URL1. The dot (.) means that the current working directory is the context2. Therefore, docker build -t myimage:1.0 . means that the image will be built from the Dockerfile and the files in the current working directory, and it will be named myimage and tagged as 1.0.

NEW QUESTION 140
An administrator thinks that a package was installed using a snap. Which of the following commands can the administrator use to verify this information?

A. snap list
B. snap find
C. snap install
D. snap try

Answer: A

Explanation:
The snap list command is used to display the installed snaps on the system1. Snaps are self-contained software packages that can be installed and updated across different Linux distributions2. The snap list command shows the name, version, revision, developer and notes of each snap1.
The snap find command is used to search for snaps in the Snap Store, which is an online repository of snaps2. The snap install command is used to install snaps from the Snap Store or from a local file2. The snap try command is used to test a snap without installing it, by mounting a directory that contains the snap files2. These commands are not useful for verifying if a package was installed using a snap.

NEW QUESTION 142
A Linux administrator needs to create a new user named user02. However, user02 must be in a different home directory, which is under /comptia/projects. Which of the following commands will accomplish this task?

A. useradd -d /comptia/projects user02
B. useradd -m /comptia/projects user02
C. useradd -b /comptia/projects user02
D. useradd -s /comptia/projects user02

Answer: A

Explanation:
The command useradd -d /comptia/projects user02 will accomplish the task of creating a new user named user02 with a different home directory.
The useradd command is a tool for creating new user accounts on Linux systems. The - d option specifies the home directory for the new user, which is the directory where the user's personal files and settings are stored. The /comptia/projects is the path of the home directory for the new user, which is different from the default location of /home/user02.
The user02 is the name of the new user. The command useradd -d /comptia/projects user02 will create a new user named user02 with a home directory under /comptia/projects. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not specify the home directory for the new user (useradd -m /comptia/projects user02 or useradd -s /comptia/projects user02) or do not use the correct option for the home directory (useradd -b /comptia/projects user02 instead of useradd -d /comptia/projects user02). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Users and Groups, page 403.

NEW QUESTION 144
Ann, a security administrator, is performing home directory audits on a Linux server. Ann issues the su Joe command and then issues the Is command. The output displays files that reside in Ann's home directory instead of Joe's. Which of the following represents the command Ann should have issued in order to list Joe's files?

A. su - Joe
B. sudo Joe
C. visudo Joe
D. pkexec joe

Answer: A

Explanation:
The su command is used to switch to another user account on Linux systems. The - option makes the shell a login shell, which means that it will read the profile and environment variables of the target user. Without this option, the shell will retain the environment variables of the original user. This can cause confusion when issuing commands that depend on these variables, such as Is, which uses the $HOME variable to determine the home directory. Therefore, Ann should have issued su - Joe to list Joe's files instead of her own. References: [How to Use su Command in Linux with Examples]

NEW QUESTION 148
A systems engineer is adding a new 1GB XFS filesystem that should be temporarily mounted under /ops/app. Which of the following is the correct list of commands to achieve this goal?

A.
```
pvcreate -L1G /dev/app
mkfs.xfs /dev/app
mount /dev/app /opt/app
```

B.

```
parted /dev/sdb --script mkpart primary xfs 1GB
mkfs.xfs /dev/sdb
mount /dev/sdb /opt/app
```

C.

```
lvs --create 1G --name app
mkfs.xfs /dev/app
mount /dev/app /opt/app
```

D.

```
lvcreate -L 1G -n app app_vq
mkfs.xfs /dev/app_vg/app
mount /dev/app_vg/app /opt/app
```

Answer: D

Explanation:
The list of commands in option D is the correct way to achieve the goal. The commands are as follows:
? fallocate -l 1G /ops/app.img creates a 1GB file named app.img under the /ops directory.
? mkfs.xfs /ops/app.img formats the file as an XFS filesystem.
? mount -o loop /ops/app.img /ops/app mounts the file as a loop device under the /ops/app directory. The other options are incorrect because they either use the wrong commands (dd or truncate instead of fallocate), the wrong options (-t or - f instead of -o), or the wrong order of arguments (/ops/app.img /ops/app instead of /ops/app /ops/app.img). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 323-324.

NEW QUESTION 149
A systems administrator made some unapproved changes prior to leaving the company. The newly hired administrator has been tasked with revealing the system to a compliant state. Which of the following commands will list and remove the correspondent packages?

A. dnf list and dnf remove last
B. dnf remove and dnf check
C. dnf info and dnf upgrade
D. dnf history and dnf history undo last

Answer: D

Explanation:
The commands that will list and remove the corresponding packages are dnf history and dnf history undo last. The dnf history command will display a list of all transactions performed by dnf, such as installing, updating, or removing packages. Each transaction has a unique ID, a date and time, an action, and a number of altered packages. The dnf history undo last command will undo the last transaction performed by dnf, meaning that it will reverse all package changes made by that transaction. For example, if the last transaction installed some packages, dnf history undo last will remove them.
The other options are not correct commands for listing and removing corresponding packages. The dnf list command will display a list of available packages in enabled repositories, but not the packages installed by dnf transactions. The dnf remove command will remove specified packages from the system, but not all packages from a specific transaction. The dnf info command will display detailed information about specified packages, but not about dnf transactions. The dnf upgrade command will upgrade all installed packages to their latest versions, but not undo any package changes. References: Handling package management history; dnf(8) - Linux manual page

NEW QUESTION 154
A Linux administrator is reviewing changes to a configuration file that includes the following section:

```
tls:
    certificates:
        - certFile: /etc/ssl/cert.cer
          keyFile: /etc/ssl/cert.key
          stores: default
        - certFile: /etc/ssl/expired.cer
          keyFile: /etc/ssl/expired.key
          stores: expired
```

The Linux administrator is trying to select the appropriate syntax formatter to correct any issues with the configuration file. Which of the following should the syntax formatter support to meet this goal?

A. Markdown
B. XML
C. YAML

D. JSON

Answer: C

Explanation:
The configuration file shown in the image is written in YAML format, so the syntax formatter should support YAML to correct any issues with the file. YAML stands for YAML Ain't Markup Language, and it is a human-readable data serialization language that uses indentation and colons to define key-value pairs. YAML supports various data types, such as scalars, sequences, mappings, anchors, aliases, and tags. The configuration file follows the rules and syntax of YAML, while the other options do not. Markdown is a lightweight markup language that uses plain text formatting to create rich text documents. XML is a markup language that uses tags to enclose elements and attributes. JSON is a data interchange format that uses curly braces to enclose objects and square brackets to enclose arrays. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 21: Automating Tasks with Ansible, page 591.

NEW QUESTION 156
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

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.

NEW QUESTION 157
A Linux administrator generated a list of users who have root-level command-line access to the Linux server to meet an audit requirement. The administrator analyzes the following /etc/passwd and /etc/sudoers files:
$ cat /etc/passwd
root:x: 0:0: :/home/root: /bin/bash lee: x: 500: 500: :/home/lee:/bin/tcsh
mallory:x: 501:501: :/root:/bin/bash
eve:x: 502: 502: /home/eve:/bin/nologin carl:x:0:503: :/home/carl:/bin/sh
bob:x: 504: 504: : /home/bob:/bin/ksh
alice:x: 505:505: :/home/alice:/bin/rsh
$ cat /etc/sudoers
Cmnd_Alias SHELLS = /bin/tcsh, /bin/sh, /bin/bash Cmnd_Alias SYSADMIN = /usr/sbin/tcpdump
ALL = (ALL) ALL
ALL = NOPASSWD: SYSADMIN
Which of the following users, in addition to the root user, should be listed in the audit report as having root-level command-line access? (Select two).

A. Carl
B. Lee
C. Mallory
D. Eve
E. Bob
F. Alice

Answer: AC

Explanation:
The users who have root-level command-line access are those who have either the same user ID (UID) as root, which is 0, or the ability to run commands as root using sudo. Based on the /etc/passwd and /etc/sudoers files, the users who meet these criteria are:
? Carl: Carl has the same UID as root, which is 0, as shown in the /etc/passwd file.
This means that Carl can log in as root and execute any command with root privileges1
? Mallory: Mallory has the ability to run commands as root using sudo, as shown in the /etc/sudoers file. The line ALL = (ALL) ALL means that any user can run any command as any other user, including root, by using sudo. Mallory can also use the root shell /bin/bash as her login shell, as shown in the /etc/passwd file2
Therefore, the correct answer is A and C. Lee, Eve, Bob, and Alice do not have root-level command-line access because they have different UIDs from root and they cannot use sudo to run commands as root. Lee can only use sudo to run the commands listed in the Cmnd_Alias SHELLS, which are /bin/tcsh, /bin/sh, and /bin/bash. Eve cannot log in at all because her login shell is /bin/nologin. Bob and Alice can only use sudo to run the command /usr/sbin/tcpdump without a password, as specified by the Cmnd_Alias SYSADMIN and the line ALL = NOPASSWD: SYSADMIN2

NEW QUESTION 162
A Linux administrator is troubleshooting the root cause of a high CPU load and average.

```
$ uptime
07:30:43 up 20 days, 3 min, 1 user, load average: 2.98, 3.62, 5.21

$ top
PID  USER PR  NI  VIRT RES   SHR  S %CPU %MEM TIME+  COMMAND
6295 user1 30 -10 5465 56465 8254 R 86.5 1.5 7:35.25 app1

$ ps -ef | grep user1
user1 6295 1 7:42:19 tty/1   06:48:29 /usr/local/bin/app1
```

Which of the following commands will permanently resolve the issue?

A. renice -n -20 6295
B. pstree -p 6295
C. iostat -cy 1 5
D. kill -9 6295

Answer: D

Explanation:
The command that will permanently resolve the issue of high CPU load and average is kill -9 6295. This command will send a SIGKILL signal to the process with the PID 6295, which is the process that is consuming 99.7% of the CPU according to the top output. The SIGKILL signal will terminate the process immediately and free up the CPU resources. The kill command is used to send signals to processes by PID or name.
The other options are not correct commands for resolving this issue. The renice -n -20 6295 command will change the priority (niceness) of the process with PID 6295 to -20, which is the highest priority possible. This will make the process more CPU-intensive, not less. The renice command is used to change the priority of running processes. The pstree - p 6295 command will show a tree of processes with PID 6295 as the root. This will not affect the CPU load or average, but only display information. The pstree command is used to display a tree of processes. The iostat -cy 1 5 command will show CPU and disk I/O statistics for 5 iterations with an interval of 1 second. This will also not affect the CPU load or average, but only display information. The iostat command is used to report CPU and I/O statistics. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Troubleshooting Linux Systems; kill(1) - Linux manual page; renice(1) - Linux manual page; pstree(1) - Linux manual page; iostat(1) - Linux manual page

NEW QUESTION 166
A Linux user reported the following error after trying to connect to the system remotely: ssh: connect to host 10.0.1.10 port 22: Resource temporarily unavailable
The Linux systems administrator executed the following commands in the Linux system while trying to diagnose this issue:

```
# netstat -an | grep 22 | grep LISTEN
tcp      0      0      0.0.0.0:22      0.0.0.0:*      LISTEN

# firewall-cmd --list-all
public (active)
  target: default
  icmp-block-inversion: no
  interfaces: eth0
  sources:
  services: dhcpv6-client
  ports:
  protocols:
  masquerade: no
        forward-ports:
  source-ports:
  icmp-blocks:
  rich rules:
```

Which of the following commands will resolve this issue?

A. firewall-cmd --zone=public --permanent --add-service=22
B. systemct1 enable firewalld; systemct1 restart firewalld
C. firewall-cmd --zone=public --permanent --add-service=ssh
D. firewall-cmd --zone=public --permanent --add-port=22/udp

Answer: C

Explanation:
The firewall-cmd --zone=public --permanent --add-service=ssh command will resolve the issue by allowing SSH connections on port 22 in the public zone of the firewalld service. This command will add the ssh service to the permanent configuration of the public zone, which means it will persist after a reboot or a reload of the firewalld service. The firewall-cmd --zone=public --permanent --add-service=22 command is invalid, as 22 is not a valid service name. The systemct1 enable firewalld; systemct1 restart firewalld command will enable and restart the firewalld service, but it will not change the firewall rules. The firewall-cmd --zone=public --permanent --add-port=22/udp command will allow UDP traffic on port 22 in the public zone, but SSH uses TCP, not UDP. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 543.

NEW QUESTION 168
A systems administrator is investigating a service that is not starting up. Given the following information:

```
root@localhost ~]# systemctl status network
network.service - LSB: Bring up/down networking
Loaded: loaded (/etc/rc.d/init.d/network; bad; vendor preset: disabled)
Active: failed (Result: exit-code) since Jan 2022-01-02 22:55:15 CST;
Docs: man:systemd-sysv-generator(8)
Process: 1083 ExecStart=/etc/rc.d/init.d/network start (code=exited, status=1/FAILURE)
Jan 02 22:55:15 localhost.localdomain network[1083]: Bringing up interface enp0s25: Error: Con...n.
Jan 02 22:55:15 localhost.localdomain network[1083]: [FAILED]
[...]
```

Which of the following systemd commands should the administrator use in order to obtain more details about the failing service?

A. systemct1 analyze network
B. systemct1 info network
C. sysctl -a network
D. journalctl -xu network

Answer: D

Explanation:
The systemd is a system and service manager for Linux systems that provides a standard way to control and monitor system services. The systemd uses various commands and tools to manage and troubleshoot system services, such as systemct1, sysctl, and journalctl. The systemct1 command is used to start, stop, enable, disable, restart, reload, status, and list system services. The sysctl command is used to configure kernel parameters at runtime. The journalctl command is used to view and filter the logs of system services.
To investigate a service that is not starting up, the administrator can use the journalctl command with the -xu option. The -x option enables verbose output that includes explanatory text and priority information. The -u option filters the output by a specific unit name, such as network.service. Therefore, the command journalctl -xu network will show detailed logs of the network service, which can help identify the cause of the failure. The statement D is correct.
The statements A, B, and C are incorrect because they do not provide more details about the failing service. The systemct1 analyze network command does not exist.
The systemct1 info network command shows basic information about the network unit, such as description, load state, active state, sub state, and main PID. The sysctl -a network command shows all kernel parameters related to network settingsR. eferences: [How to Use Systemd to Manage System Services]

NEW QUESTION 170
.....

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

XK0-005 Practice Exam Features:

* XK0-005 Questions and Answers Updated Frequently

* XK0-005 Practice Questions Verified by Expert Senior Certified Staff

* XK0-005 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* XK0-005 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The XK0-005 Practice Test Here