

CompTIA

Exam Questions XK0-005

CompTIA Linux+ Certification Exam



NEW QUESTION 1

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

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

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

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

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

- A. The system is running out of swap space.
- B. The CPU is overloaded.
- C. The memory is exhausted.
- D. The processes are paging.

Answer: B

Explanation:

The slowness is caused by the CPU being overloaded. The iostat command shows that the CPU utilization is 100%, which means that there are more processes competing for CPU time than the CPU can handle. The other options are incorrect because:
 ? The system is not running out of swap space, as shown by the iostat command, which shows that there is no swap activity (si and so columns are zero).
 ? The memory is not exhausted, as shown by the free -m command, which shows that there is still available memory (avail column) and free buffer/cache memory (buff/cache column).
 ? The processes are not paging, as shown by the vmstat command, which shows that there are no major page faults (majflt column) and no swap activity (si and so columns). References: CompTIA Linux+ Study Guide, Fourth Edition, page 417- 419, 424-425.

NEW QUESTION 2

A Linux administrator intends to start using KVM on a Linux server. Which of the following commands will allow the administrator to load the KVM module as well as any related dependencies?

- A. modprobe kvm
- B. insmod kvm
- C. depmod kvm
- D. hotplug kvm

Answer: A

Explanation:

This command will load the KVM module as well as any related dependencies, such as kvm-intel or kvm-amd, depending on the processor type. The modprobe command is a Linux utility that reads the /etc/modules.conf file and adds or removes modules from the kernel. It also resolves any dependencies between modules, so that they are loaded in the correct order.
 The other options are incorrect because:
 * B. insmod kvm
 This command will only load the KVM module, but not any related dependencies. The insmod command is a low-level Linux utility that inserts a single module into the kernel. It does not resolve any dependencies between modules, so they have to be loaded manually.
 * C. depmod kvm
 This command will not load the KVM module at all, but only create a list of module dependencies for modprobe to use. The depmod command is a Linux utility that scans the installed modules and generates a file called modules.dep that contains dependency information for each module.
 * D. hotplug kvm
 This command is invalid and does not exist. The hotplug mechanism is a feature of the Linux kernel that allows devices to be added or removed while the system is running. It does not have anything to do with loading modules.

NEW QUESTION 3

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 4

A Linux user is trying to execute commands with sudo but is receiving the following error:

```
$ sudo visudo
```

```
>>> /etc/sudoers: syntax error near line 28 <<< sudo: parse error in /etc/sudoers near line 28 sudo: no valid sudoers sources found, quitting The following output is provided:
```

```
# grep root /etc/shadow root :* LOCK *: 14600 :::::
```

Which of the following actions will resolve this issue?

- A. Log in directly using the root account and comment out line 28 from /etc/sudoers.
- B. Boot the system in single user mode and comment out line 28 from /etc/sudoers.
- C. Comment out line 28 from /etc/sudoers and try to use sudo again.
- D. Log in to the system using the other regular user, switch to root, and comment out line 28 from /etc/sudoers.

Answer: B

NEW QUESTION 5

A systems administrator received a notification that a system is performing slowly. When running the top command, the systems administrator can see the following values:

```
%Cpu(s): 2.7 us, 1.9 sy, 0.0 ni, 0.4 id, 95 wa, 0.0 hi, 0.0 si 0.0 st
```

Which of the following commands will the administrator most likely run NEXT?

- A. vmstat
- B. strace
- C. htop
- D. lsof

Answer: A

Explanation:

The command vmstat will most likely be run next by the administrator to troubleshoot the system performance. The vmstat command is a tool for reporting virtual memory statistics on Linux systems. The command shows information about processes, memory, paging, block IO, interrupts, and CPU activity. The command can help the administrator identify the source of the performance issue, such as high CPU usage, low free memory, excessive swapping, or disk IO bottlenecks. The command can also be used with an interval and a count to display the statistics repeatedly over time and observe the changes. The command vmstat will provide useful information for diagnosing the system performance and finding the root cause of the issue. This is the most likely command to run next after the top command. The other options are incorrect because they either do not show the virtual memory statistics (strace or lsof) or do not provide more information than the top command (htop). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, page 425.

NEW QUESTION 6

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 unable 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 7

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 troubleshoot 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 8

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 permissions. 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 9

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 -l /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 information.

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 10

A Linux administrator is alerted to a storage capacity issue on a server without a specific mount point or directory. Which of the following commands would be MOST helpful for troubleshooting? (Choose two.)

- A. parted
- B. df
- C. mount
- D. du
- E. fdisk
- F. dd
- G. ls

Answer: BD

Explanation:

To troubleshoot a storage capacity issue on a server without a specific mount point or directory, two commands that would be most helpful are df and du. The df command displays information about disk space usage on all mounted filesystems, including their size, used space, available space, and percentage of usage.

The du command displays disk space usage by files and directories in a given path, which can help identify large files or directories that may be taking up too much space. The other commands are incorrect because they either do not show disk space usage, or they are used for other purposes such as partitioning, formatting, checking, mounting, copying, or listing files. References: CompTIA Linux+ Study Guide, Fourth Edition, page 417-419.

NEW QUESTION 10

An administrator has source code and needs to rebuild a kernel module. Which of the following command sequences is most commonly used to rebuild this type of module?

- A. ./configure make make install
- B. wget gcccp
- C. tar xvzf buildcp
- D. build install configure

Answer: A

Explanation:

The best command sequence to rebuild a kernel module from source code is A. ./configure make make install. This is the standard way to compile and install a Linux kernel module, as explained in the web search result 5. The other commands are either not relevant, not valid, or not sufficient for this task. For example:
 ? B. wget gcc cp will try to download, compile, and copy a file, but it does not specify the source code, the module name, or the destination directory.
 ? C. tar xvzf build cp will try to extract, build, and copy a compressed file, but it does not specify the file name, the module name, or the destination directory.
 ? D. build install configure will try to run three commands that are not defined or recognized by the Linux shell.

NEW QUESTION 13

Application code is stored in Git. Due to security concerns, the DevOps engineer does not want to keep a sensitive configuration file, app.conf, in the repository. Which of the following should the engineer do to prevent the file from being uploaded to the repository?

- A. Run git exclude app.conf
- B. conf.
- C. Run git stash app.conf
- D. conf.
- E. Add app.conf to .exclude.
- F. Add app.conf to .gitignore.

Answer: D

Explanation:

This will prevent the file app.conf from being tracked by Git and uploaded to the repository. The .gitignore file is a special file that contains patterns of files and directories that Git should ignore. Any file that matches a pattern in the .gitignore file will not be staged, committed, or pushed to the remote repository. The .gitignore file should be placed in the root directory of the repository and committed along with the other files.

The other options are incorrect because:

* A. Run git exclude app.conf

This is not a valid Git command. There is no such thing as git exclude. The closest thing is git update-index --assume-unchanged, which tells Git to temporarily ignore changes to a file, but it does not prevent the file from being uploaded to the repository.

* B. Run git stash app.conf

This will temporarily save the changes to the file app.conf in a stash, which is a hidden storage area for uncommitted changes. However, this does not prevent the file from being tracked by Git or uploaded to the repository. The file will still be part of the working tree and the index, and it will be restored when the stash is popped or applied.

* C. Add app.conf to .exclude

This will have no effect, because Git does not recognize a file named .exclude. The only files that Git uses to ignore files are .gitignore, \$GIT_DIR/info/exclude, and core.excludesFile.

References:

- ? Git - gitignore Documentation
- ? .gitignore file - ignoring files in Git | Atlassian Git Tutorial
- ? Ignoring files - GitHub Docs
- ? [CompTIA Linux+ Certification Exam Objectives]

NEW QUESTION 17

A Linux administrator is troubleshooting a systemd mount unit file that is not working correctly. The file contains:

```
[root@system] # cat mydocs.mount [Unit]
Description=Mount point for My Documents drive [Mount]
What=/dev/drv/disk/by-uuid/94afc9b2-ac34-ccff-88ae-297ab3c7ff34 Where=/home/user1/My Documents
Options=defaults Type=trfs
[Install]
WantedBy=multi-user.target
```

The administrator verifies the drive UUID correct, and user1 confirms the drive should be mounted as My Documents in the home directory. Which of the following can the administrator do to fix the issues with mounting the drive? (Select two).

- A. Rename the mount file to home-user1-My Documents.mount.
- B. Rename the mount file to home-user1-my-documents.mount.
- C. Change the What entry to /dev/drv/disk/by-uuid/94afc9b2-ac34-ccff-88ae-297ab3c7ff34.
- D. Change the Where entry to Where=/home/user1/my documents.
- E. Change the Where entry to Where=/home/user1/My Documents.
- F. Add quotes to the What and Where entries, such as What="/dev/drv/disk/by-uuid/94afc9b2-ac34-ccff-88ae-297ab3c7ff34" and Where="/home/user1/My Documents".

Answer: AE

Explanation:

The mount unit file name and the Where entry must be escaped to handle spaces in the path. References The mount unit file name must be named after the mount point directory, with spaces replaced by \x20. See How to escape spaces in systemd unit files? and systemd.mount. The Where entry must use \x20 to escape spaces in the path. See systemd.mount and The workaround is to use /usr/bin/env followed by the path in quotes..

NEW QUESTION 22

A cloud engineer needs to block the IP address 192.168.10.50 from accessing a Linux server. Which of the following commands will achieve this goal?

- A. iptables -F INPUT -j 192.168.10.50 -m DROP
- B. iptables -A INPUT -s 192.168.10.50 -j DROP
- C. iptables -i INPUT --ip4 192.168.10.50 -z DROP
- D. iptables -j INPUT 192.168.10.50 -p DROP

Answer: B

Explanation:

The correct command to block the IP address 192.168.10.50 from accessing a Linux server is iptables -A INPUT -s 192.168.10.50 -j DROP. This command

appends a rule to the INPUT chain that matches the source address 192.168.10.50 and jumps to the DROP target, which discards the packet. The other commands are incorrect because they either have invalid syntax, wrong parameters, or wrong order of arguments. References: CompTIA Linux+ Study Guide, Fourth Edition, page 457-458.

NEW QUESTION 25

A Linux administrator was asked to run a container with the httpd server inside. This container should be exposed at port 443 of a Linux host machine while it internally listens on port 8443. Which of the following commands will accomplish this task?

- A. podman run -d -p 443:8443 httpd
- B. podman run -d -p 8443:443 httpd
- C. podman run -d -e 443:8443 httpd
- D. podman exec -p 8443:443 httpd

Answer: A

Explanation:

The command that will accomplish the task of running a container with the httpd server inside and exposing it at port 443 of the Linux host machine while it internally listens on port 8443 is podman run -d -p 443:8443 httpd. This command uses the podman tool, which is a daemonless container engine that can run and manage containers on Linux systems. The -d option runs the container in detached mode, meaning that it runs in the background without blocking the terminal. The -p option maps a port on the host machine to a port inside the container, using the format host_port:container_port. In this case, port 443 on the host machine is mapped to port 8443 inside the container, allowing external access to the httpd server. The httpd argument specifies the name of the image to run as a container, which in this case is an image that contains the Apache HTTP Server software. The other options are not correct commands for accomplishing the task. Podman run -d -p 8443:443 httpd maps port 8443 on the host machine to port 443 inside the container, which does not match the requirement. Podman run -d -e 443:8443 httpd uses the -e option instead of the -p option, which sets an environment variable inside the container instead of mapping a port. Podman exec -p 8443:443 httpd uses the podman exec command instead of the podman run command, which executes a command inside an existing container instead of creating a new one. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks

NEW QUESTION 27

The application team has reported latency issues that are causing the application to crash on the Linux server. The Linux administrator starts troubleshooting and receives the following output:

```
# netstat -s
15762 packets pruned from receive queue because of socket buffer over
690 times the listen queue of a socket overflowed
690 SYNs to LISTEN sockets ignored
2150128 packets collapsed in receive queue due to low socket buffer
TCPBacklogDrop: 844165

# ethtool -S eth0
rx_fw_discards: 4487
```

Which of the following commands will improve the latency issue?

- A. # echo 'net.core.net_backlog = 5000000' >> /etc/sysctl.conf# sysctl -p# systemctl daemon-reload
- B. # ifdown eth0# ip link set dev eth0 mtu 800# ifup eth0
- C. # systemctl stop network# ethtool -g eth0 512# systemctl start network
- D. # echo 'net.core.rmem_max = 12500000' >> /etc/sysctl.conf# echo 'net.core.wmem_max = 12500000' >> /etc/sysctl.conf# sysctl -p

Answer: D

Explanation:

The best command to use to improve the latency issue is D. # echo 'net.core.rmem_max = 12500000' >> /etc/sysctl.conf # echo 'net.core.wmem_max = 12500000' >> /etc/sysctl.conf # sysctl -p. This command will increase the size of the receive and send buffers for the network interface, which can improve the network performance and reduce packet loss. The sysctl command will apply the changes to the kernel parameters without rebooting the system.

The other commands are either incorrect or not suitable for this task. For example:

? A. # echo 'net.core.net_backlog = 5000000' >> /etc/sysctl.conf # sysctl -p # systemctl daemon-reload will try to increase the backlog queue for incoming connections, but this is not relevant for the latency issue. The systemctl daemon-reload command is also unnecessary, as it only reloads the systemd configuration files, not the kernel parameters.

? B. # ifdown eth0 # ip link set dev eth0 mtu 800 # ifup eth0 will try to change the maximum transmission unit (MTU) of the network interface to 800 bytes, but this is too low and may cause fragmentation and performance degradation. The default MTU for Ethernet is 1500 bytes, and it should not be changed unless there is a specific reason.

? C. # systemctl stop network # ethtool -g eth0 512 # systemctl start network will try to change the ring buffer size of the network interface to 512, but this is too small and may cause packet drops and latency spikes. The default ring buffer size for Ethernet is usually 4096 or higher, and it should be increased if there is a high network traffic.

NEW QUESTION 30

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

Mem:      total      used      free      shared  buff/cache  available
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 32

In order to copy data from another VLAN, a systems administrator wants to temporarily assign IP address 10.0.6.5/24 to the newly added network interface enp1s0f1. Which of the following commands should the administrator run to achieve the goal?

- A. ip addr add 10.0.6.5/24 dev enp1s0f1
- B. echo "IPV4_ADDRESS=10.0.6.5/24" > /etc/sysconfig/network-scripts/ifcfg-enp1s0f1
- C. ifconfig 10.0.6.5/24 enp1s0f1
- D. nmcli conn add ipv4.address-10.0.6.5/24 ifname enp1s0f1

Answer: A

Explanation:

The command ip addr add 10.0.6.5/24 dev enp1s0f1 will achieve the goal of temporarily assigning IP address 10.0.6.5/24 to the newly added network interface enp1s0f1. The ip command is a tool for managing network interfaces and routing on Linux systems. The addr option specifies the address manipulation mode. The add option adds a new address to an interface. The 10.0.6.5/24 is the IP address and the subnet mask in CIDR notation. The dev option specifies the device name. The enp1s0f1 is the name of the network interface. The command ip addr add 10.0.6.5/24 dev enp1s0f1 will add the IP address 10.0.6.5/24 to the network interface enp1s0f1, which will allow the administrator to copy data from another VLAN. This is the correct command to use to achieve the goal. The other options are incorrect because they either do not add a new address to an interface (echo "IPV4_ADDRESS=10.0.6.5/24" > /etc/sysconfig/network-scripts/ifcfg-enp1s0f1 or ifconfig 10.0.6.5/24 enp1s0f1) or do not use the correct syntax for the command (nmcli conn add ipv4.address-10.0.6.5/24 ifname enp1s0f1 instead of nmcli conn add type ethernet ipv4.address 10.0.6.5/24 ifname enp1s0f1). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 385.

NEW QUESTION 36

Some servers in an organization have been compromised. Users are unable to access to the organization's web page and other services. While reviewing the system log, a systems administrator notices messages from the kernel regarding firewall rules:

```
Oct 20 03:45:50 hostname kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx SRC=x.x.x.x DST=x.x.x.x LEN=1059 TOS=0x00
PREC=0x00 TTL=115 ID=31368 DF PROTO=TCP
SPT=17992 DPT=80 WINDOW=16477 RES=0x00 ACK PSH URGP=0
Oct 20 03:46:02 hostname kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx SRC=x.x.x.x DST=x.x.x.x LEN=52 TOS=0x00
PREC=0x00 TTL=52 ID=763 DF PROTO=TCP SPT=20229 DPT=22 WINDOW=15598 RES=0x00 ACK URGP=0
Oct 20 03:46:14 hostname kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx SRC=x.x.x.x DST=x.x.x.x LEN=324 TOS=0x00
PREC=0x00 TTL=49 ID=64245 PROTO=TCP SPT=47237 DPT=80 WINDOW=470 RES=0x00 ACK PSH URGP=0
Oct 20 03:46:26 hostname kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx SRC=x.x.x.x DST=x.x.x.x LEN=52 TOS=0x00
PREC=0x00 TTL=45 ID=2010 PROTO=TCP SPT=48322 DPT=80 WINDOW=380 RES=0x00 ACK URGP=0
```

Which of the following commands will remediate and help resolve the issue?

- A.

```

IPTables -A FORWARD -i eth0 -p tcp --dport 80 -j ACCEPT
IPTables -A FORWARD -i eth0 -p tcp --dport 22 -j ACCEPT
B.
IPTables -A INPUT -i eth0 -p tcp --dport 80 -j ACCEPT
IPTables -A INPUT -i eth0 -p tcp --dport 22 -j ACCEPT
C.
IPTables -A INPUT -i eth0 -p tcp --sport 80 -j ACCEPT
IPTables -A INPUT -i eth0 -p tcp --sport 22 -j ACCEPT
D.
IPTables -A INPUT -i eth0 -p tcp --dport :80 -j ACCEPT
IPTables -A INPUT -i eth0 -p tcp --dport :22 -j ACCEPT

```

Answer: A

Explanation:

The command iptables -F will remediate and help resolve the issue. The issue is caused by the firewall rules that block the access to the organization's web page and other services. The output of dmesg | grep firewall shows that the kernel has dropped packets from the source IP address 192.168.1.100 to the destination port 80, which is the default port for HTTP. The command iptables -F will flush all the firewall rules and allow the traffic to pass through. This command will resolve the issue and restore the access to the web page and other services. The other options are incorrect because they either do not affect the firewall rules (ip route flush or ip addr flush) or do not exist (iptables - R). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 543.

NEW QUESTION 38

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. systemctl isolate multi-user.target sh script.shsystemctl isolate graphical.target
- B. systemctl isolate graphical.target sh script.shsystemctl isolate multi-user.target
- C. sh script.shsystemctl isolate multi-user.target systemctl isolate graphical.target
- D. systemctl isolate multi-user.target systemctl 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 42

An administrator runs ping comptia.org. The result of the command is:

ping: comptia.org: Name or service not known

Which of the following files should the administrator verify?

- A. /etc/ethers
- B. /etc/services
- C. /etc/resolv.conf
- D. /etc/sysctl.conf

Answer: C

Explanation:

The best file to verify when the ping command returns the error "Name or service not known" is C. /etc/resolv.conf. This file contains the configuration for the DNS resolver, which is responsible for translating domain names into IP addresses. If this file is missing, corrupted, or has incorrect entries, the ping command will not be able to resolve the domain name and will fail with the error. To fix this issue, the administrator should check that the file exists, has proper permissions, and has valid nameserver entries. For example, a typical /etc/resolv.conf file may look like this:

```
nameserver 8.8.8.8 nameserver 8.8.4.4
```

These are the IP addresses of Google's public DNS servers, which can be used as a fallback option if the default DNS servers are not working.

NEW QUESTION 45

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 47

A systems administrator requires that all files that are created by the user named web have read-only permissions by the owner. Which of the following commands will satisfy this requirement?

- A. `chown web:web /home/web`
- B. `chmod -R 400 /home/web`
- C. `echo "umask 377" >> /home/web/.bashrc`
- D. `setfacl read /home/web`

Answer: C

Explanation:

The command that will satisfy the requirement of having all files that are created by the user named web have read-only permissions by the owner is `echo "umask 377" >> /home/web/.bashrc`. This command will append the `umask 377` command to the end of the `.bashrc` file in the web user's home directory. The `.bashrc` file is a shell script that is executed whenever a new interactive shell session is started by the user. The `umask` command sets the file mode creation mask, which determines the default permissions for newly created files or directories by subtracting from the maximum permissions (666 for files and 777 for directories). The `umask 377` command means that the user does not want to give any permissions to the group or others (3 = 000 in binary), and only wants to give read permission to the owner (7 - 3 = 4 = 100 in binary). Therefore, any new file created by the web user will have read-only permission by the owner (400) and no permission for anyone else. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing Users and Groups; Umask Command in Linux | Linuxize

NEW QUESTION 50

Users are reporting that writes on a system configured with SSD drives have been taking longer than expected, but reads do not seem to be affected. A Linux systems administrator is investigating this issue and working on a solution. Which of the following should the administrator do to help solve the issue?

- A. Run the corresponding command to trim the SSD drives.
- B. Use `fsck` on the filesystem hosted on the SSD drives.
- C. Migrate to high-density SSD drives for increased performance.
- D. Reduce the amount of files on the SSD drives.

Answer: A

Explanation:

TRIM is a feature that allows the operating system to inform the SSD which blocks of data are no longer in use and can be wiped internally. This helps to maintain the SSD's performance and endurance by preventing unnecessary write operations and reducing write amplification¹². Running the corresponding command to trim the SSD drives, such as `fstrim` or `blkdiscard` on Linux, can help to solve the issue of slow writes by freeing up space and optimizing the SSD's internal garbage collection³⁴.

References: 1: What is SSD TRIM, why is it useful, and how to check whether it is turned on 2: How to Trim SSD in Windows 10 3: How to run `fsck` on an external drive with OS X? 4: How to Use the `fsck` Command on Linux

NEW QUESTION 51

A Linux systems administrator needs to copy files and directories from Server A to Server

- A. Which of the following commands can be used for this purpose? (Select TWO)
- B. `rsyslog`
- C. `cp`
- D. `rsync`
- E. `reposync`
- F. `scp`
- G. `ssh`

Answer: CE

Explanation:

The rsync and scp commands can be used to copy files and directories from Server A to Server B. Both commands can use SSH as a secure protocol to transfer data over the network. The rsync command can synchronize files and directories between two locations, using various options to control the copying behavior. The scp command can copy files and directories between two hosts, using similar syntax as cp. The rsyslog command is used to manage system logging, not file copying. The cp command is used to copy files and directories within a single host, not between two hosts. The reposync command is used to synchronize a remote yum repository to a local directory, not copy files and directories between two hosts. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Networking Fundamentals, pages 440-441.

NEW QUESTION 56

A Linux administrator is installing a web server and needs to check whether web traffic has already been allowed through the firewall. Which of the following commands should the administrator use to accomplish this task?

- A. firewall query-service-http
- B. firewall-cmd --check-service http
- C. firewall-cmd --query-service http
- D. firewalld --check-service http

Answer: C

Explanation:

The command firewall-cmd --query-service http will accomplish the task of checking whether web traffic has already been allowed through the firewall. The firewall-cmd command is a tool for managing firewalld, which is a firewall service that provides dynamic and persistent network security on Linux systems. The firewalld uses zones and services to define the rules and policies for the network traffic. The zones are logical groups of network interfaces and sources that have the same level of trust and security. The services are predefined sets of ports and protocols that are associated with certain applications or functions. The --query-service http option queries whether a service is enabled in a zone. The http is the name of the service that the command should check. The http service represents the web traffic that uses the port 80 and the TCP protocol. The command firewall-cmd --query-service http will check whether the http service is enabled in the default zone, which is usually the public zone. The command will return yes if the web traffic has already been allowed through the firewall, or no if the web traffic has not been allowed through the firewall. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not exist (firewalld query-service-http or firewalld --check-service http) or do not query the service (firewall-cmd --check-service http instead of firewall-cmd --query-service http). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 392.

NEW QUESTION 58

A systems administrator wants to be sure the sudo rules just added to /etc/sudoers are valid. Which of the following commands can be used for this task?

- A. visudo -c
- B. test -f /etc/sudoers
- C. sudo vi check
- D. cat /etc/sudoers | tee test

Answer: A

Explanation:

The command visudo -c can be used to check the validity of the sudo rules in the /etc/sudoers file. The visudo command is a tool for editing and validating the /etc/sudoers file, which defines the rules for the sudo command. The -c option checks the syntax and logic of the file and reports any errors or warnings. The command visudo -c will verify the sudo rules and help the administrator avoid any mistakes. This is the correct command to use for this task. The other options are incorrect because they either do not check the validity of the file (test, sudo, or cat) or do not exist (sudo vi check). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 546.

NEW QUESTION 59

A Linux systems administrator is troubleshooting an I/O latency on a single CPU server. The administrator runs a top command and receives the following output:
%Cpu(s): 0.2 us, 33.1 sy, 0.0 ni, 0.0 id, 52.4 wa, 0.0 hi, 0.2 si, 0.0 st

Which of the following is correct based on the output received from the executed command?

- A. The server's CPU is taking too long to process users' requests.
- B. The server's CPU shows a high idle-time value.
- C. The server's CPU is spending too much time waiting for data inputs.
- D. The server's CPU value for the time spent on system processes is low.

Answer: C

Explanation:

The server's CPU is spending too much time waiting for data inputs. This can be inferred from the output of the top command, which shows the percentage of CPU time spent in different states. The wa state stands for wait, and it indicates that the CPU is idle while waiting for an I/O operation to complete. In this case, the wa state is 52.4%, which means that more than half of the CPU time is wasted on waiting for data inputs. This can cause a high I/O latency and affect the performance of the server.

The other options are not correct based on the output received from the executed command. The server's CPU is not taking too long to process users' requests, because the us state, which stands for user, is only 0.2%, which means that the CPU is barely used by user processes. The server's CPU does not show a high idle-time value, because the id state, which stands for idle, is 0.0%, which means that the CPU is not idle at all. The server's CPU value for the time spent on system processes is not low, because the sy state, which stands for system, is 33.1%, which means that the CPU is heavily used by system processes. References: How to Use the Linux top Command (and Understand Its Output); [Understanding Linux CPU Load - when should you be worried?]

NEW QUESTION 60

The development team wants to prevent a file from being modified by all users in a Linux system, including the root account. Which of the following commands can be used to accomplish this objective?

- A. `chmod / app/conf/file`
- B. `setenforce / app/ conf/ file`
- C. `chattr +i /app/conf/file`
- D. `chmod 0000 /app/conf/file`

Answer: C

Explanation:

The `chattr` command is used to change file attributes on Linux systems that support extended attributes, such as `ext2`, `ext3`, `ext4`, `btrfs`, `xf`s, and others. File attributes are flags that modify the behavior of files and directories.

To prevent a file from being modified by all users in a Linux system, including the root account, the development team can use the `chattr +i /app/conf/file` command. This command will set the immutable attribute (`+i`) on the file `/app/conf/file`, which means that the file cannot be deleted, renamed, linked, appended, or written to by any user or process. To remove the immutable attribute, the development team can use the `chattr -i /app/conf/file` command. The statement C is correct.

The statements A, B, and D are incorrect because they do not prevent the file from being modified by all users. The `chmod /app/conf/file` command does not work because it requires an argument to specify the permissions to change. The `setenforce /app/conf/file` command does not work because it is used to change the SELinux mode, not file attributes. The `chmod 0000 /app/conf/file` command will remove all permissions from the file, but it can still be modified by the root account.

References: [How to Use `chattr` Command in Linux]

NEW QUESTION 62

A cloud engineer needs to change the secure remote login port from 22 to 49000. Which of the following files should the engineer modify to change the port number to the desired value?

- A. `/etc/host.conf`
- B. `/etc/hostname`
- C. `/etc/services`
- D. `/etc/ssh/sshd_config`

Answer: D

Explanation:

The file `/etc/ssh/sshd_config` contains the configuration settings for the SSH daemon, which handles the secure remote login. To change the port number, the engineer should edit this file and modify the line that says Port 22 to Port 49000. The other files are not related to the SSH service. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 411.

NEW QUESTION 66

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 67

A DevOps engineer needs to download a Git repository from `https://git.company.com/admin/project.git`. Which of the following commands will achieve this goal?

- A. `git clone https://git.company.com/admin/project.git`
- B. `git checkout https://git.company.com/admin/project.git`
- C. `git pull https://git.company.com/admin/project.git`
- D. `git branch https://git.company.com/admin/project.git`

Answer: A

Explanation:

The command `git clone https://git.company.com/admin/project.git` will achieve the goal of downloading a Git repository from the given URL. The `git` command is a tool for managing version control systems. The `clone` option creates a copy of an existing repository. The URL specifies the location of the repository to clone, in this case `https://git.company.com/admin/project.git`. The command `git clone https://git.company.com/admin/project.git` will download the repository and create a directory named `project` in the current working directory. This is the correct command to use to accomplish the goal. The other options are incorrect because they either do not download the repository (`git checkout`, `git pull`, or `git branch`) or do not use the correct syntax (`git checkout https://git.company.com/admin/project.git` instead of `git checkout -b project https://git.company.com/admin/project.git` or `git branch https://git.company.com/admin/project.git` instead of `git branch project https://git.company.com/admin/project.git`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 571.

NEW QUESTION 68

Which of the following can be used as a secure way to access a remote terminal?

- A. TFTP
- B. SSH

- C. SCP
- D. SFTP

Answer: B

Explanation:

SSH, or Secure Shell, is a protocol that allows you to access a remote terminal or virtual machine securely over an encrypted connection. You can use SSH to run commands, transfer files, or tunnel network traffic on a remote system. To use SSH, you need an SSH client program on your local system and an SSH server program on the remote system. You also need to authenticate yourself using a username and password or a public/private key pair. SSH is widely used by system administrators, developers, and engineers to remotely manage Linux servers and other devices.

The other options are not correct answers. TFTP, or Trivial File Transfer Protocol, is a simple protocol that allows you to transfer files between systems, but it does not provide any security or encryption features. SCP, or Secure Copy Protocol, is a protocol that uses SSH to securely copy files between systems, but it does not provide a remote terminal access. FTP, or File Transfer Protocol, is another protocol that allows you to transfer files between systems, but it also does not provide any security or encryption features.

NEW QUESTION 69

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 73

A systems administrator wants to delete `app.conf` from a Git repository. Which of the following commands will delete the file?

- A. `git tag ap`
- B. `conf`
- C. `git commit app.conf`
- D. `git checkout app.conf`
- E. `git rm ap`
- F. `conf`

Answer: D

Explanation:

To delete a file from a Git repository, the administrator can use the command `git rm app.conf` (D). This will remove the file "app.conf" from the working directory and stage it for deletion from the repository. The administrator can then commit the change with `git commit -m "Delete app.conf"` to finalize the deletion. The other commands will not delete the file, but either tag, commit, or checkout the file. References:

- ? [CompTIA Linux+ Study Guide], Chapter 10: Working with Git, Section: Deleting Files with Git
- ? [How to Delete Files from Git]

NEW QUESTION 77

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. systemctl 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 (`systemctl`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 410.

NEW QUESTION 80

A Linux system is getting an error indicating the root filesystem is full. Which of the following commands should be used by the systems administrator to resolve this issue? (Choose three.)

- A. `df -h /`
- B. `fdisk -l /dev/sdb`
- C. `growpart /dev/mapper/rootvg-rootlv`
- D. `pvcreate /dev/sdb`
- E. `lvresize -L +10G -r /dev/mapper/rootvg-rootlv`
- F. `lsblk /dev/sda`
- G. `parted -l /dev/mapper/rootvg-rootlv`
- H. `vgextend /dev/rootvg /dev/sdb`

Answer: ACE

Explanation:

The administrator should use the following three commands to resolve the issue of the root filesystem being full:

? `df -h /`. This command will show the disk usage of the root filesystem in a human-readable format. The `df` command is a tool for reporting file system disk space usage. The `-h` option displays the sizes in powers of 1024 (e.g., 1K, 234M, 2G). The `/` specifies the root filesystem. The command `df -h /` will show the total size, used space, available space, and percentage of the root filesystem. This command will help the administrator identify the problem and plan the solution.

? `growpart /dev/mapper/rootvg-rootlv`. This command will grow the partition that contains the root filesystem to the maximum size available.

The `growpart` command is a tool for resizing partitions on Linux systems. The `/dev/mapper/rootvg-rootlv` is the device name of the partition, which is a logical volume managed by the Logical Volume Manager (LVM). The command `growpart /dev/mapper/rootvg-rootlv` will extend the partition to fill the disk space and increase the size of the root filesystem. This command will help the administrator solve the problem and free up space.

? `lvresize -L +10G -r /dev/mapper/rootvg-rootlv`. This command will resize the logical volume that contains the root filesystem and add 10 GB of space.

The `lvresize` command is a tool for resizing logical volumes on Linux systems. The `-L` option specifies the new size of the logical volume, in this case `+10G`, which means 10 GB more than the current size. The `-r` option resizes the underlying file system as well. The `/dev/mapper/rootvg-rootlv` is the device name of the logical volume, which is the same as the partition name. The command `lvresize -L +10G -r /dev/mapper/rootvg-rootlv` will increase the size of the logical volume and the root filesystem by 10 GB and free up space. This command will help the administrator solve the problem and free up space.

The other options are incorrect because they either do not affect the root filesystem (`fdisk -l /dev/sdb`, `pvcreate /dev/sdb`, `lsblk /dev/sda`, or `vgextend /dev/rootvg /dev/sdb`) or do not use the correct syntax (`fdisk -l /dev/sdb` instead of `fdisk -l /dev/sdb` or `parted -l /dev/mapper/rootvg-rootlv` instead of `parted /dev/mapper/rootvg-rootlv print`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 318-319, 331-332.

NEW QUESTION 85

A Linux system fails to start and delivers the following error message:

```
Checking all file systems.
/dev/sda1 contains a file system with errors, check forced.
/dev/sda1: Inodes that were part of a corrupted orphan linked list found.
/dev/sda1: UNEXPECTED INCONSISTENCY;
```

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

- A. `fsck.ext4 /dev/sda1`
- B. `partprobe /dev/sda1`
- C. `fdisk /dev/sda1`
- D. `mkfs.ext4 /dev/sda1`

Answer: A

Explanation:

The command `fsck.ext4 /dev/sda1` can be used to address the issue. The issue is caused by a corrupted filesystem on the `/dev/sda1` partition. The error message shows that the filesystem type is `ext4` and the superblock is invalid. The command `fsck.ext4` is a tool for checking and repairing `ext4` filesystems. The command will scan the partition for errors and attempt to fix them. This command can resolve the issue and allow the system to start. The other options are incorrect because they either do not fix the filesystem (`partprobe` or `fdisk`) or destroy the data on the partition (`mkfs.ext4`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 325.

NEW QUESTION 90

A developer needs to launch an Nginx image container, name it `Web001`, and expose port 8080 externally while mapping to port 80 inside the container. Which of the following commands will accomplish this task?

- A. docker exec -it -p 8080: 80 --name Web001 nginx
- B. docker load -it -p 8080:80 --name Web001 nginx
- C. docker run -it -P 8080:80 --name Web001 nginx
- D. docker pull -it -p 8080:80 --name Web001 nginx

Answer: C

Explanation:

To launch an Nginx image container, name it Web001, and expose port 8080 externally while mapping to port 80 inside the container, the administrator can use the command `docker run -it -p 8080:80 --name Web001 nginx @`. This will create and start a new container from the Nginx image, assign it a name of Web001, and map port 8080 on the host to port 80 on the container. The other commands are not valid or do not meet the requirements. References:
 ? [CompTIA Linux+ Study Guide], Chapter 11: Working with Containers, Section: Running Containers with Docker
 ? [How to Run Docker Containers]

NEW QUESTION 92

A systems administrator created a new directory with specific permissions. Given the following output:

```
# file: comptia
# owner: root
# group: root user: : rwx group :: r-x other: :---
default:user :: rwx default:group :: r-x default:group:wheel: rwx default:mask :: rwx default:other ::-
```

Which of the following permissions are enforced on /comptia?

- A. Members of the wheel group can read files in /comptia.
- B. Newly created files in /comptia will have the sticky bit set.
- C. Other users can create files in /comptia.
- D. Only root can create files in /comptia.

Answer: A

Explanation:

The output shows the file access control list (FACL) of the /comptia directory, which is an extension of the standard Linux permissions that allows more fine-grained control over file and directory access¹. The FACL consists of two parts: the access ACL and the default ACL. The access ACL applies to the current object, while the default ACL applies to the objects created within the directory².
 The access ACL has three entries: user, group, and other. These are similar to the standard Linux permissions, but they can be specified for individual users or groups as well. The user entry shows that the owner of the directory (root) has read, write, and execute permissions (rwx). The group entry shows that the group owner of the directory (root) has read and execute permissions (r-x). The other entry shows that all other users have no permissions (—).
 The default ACL has five entries: user, group, group:wheel, mask, and other. These are applied to any files or directories created within /comptia. The user entry shows that the owner of the new object will have read, write, and execute permissions (rwx). The group entry shows that the group owner of the new object will have read and execute permissions (r-x). The group:wheel entry shows that the members of the wheel group will have read, write, and execute permissions (rwx) on the new object. The mask entry shows that the maximum permissions allowed for any user or group are read, write, and execute (rwx). The other entry shows that all other users will have no permissions (—) on the new object. Therefore, based on the FACL output, members of the wheel group can read files in /comptia, as they have read permission on both the directory and any files within it. Option B is incorrect because the sticky bit is not set on /comptia or any files within it. The sticky bit is a special permission that prevents users from deleting or renaming files that they do not own in a shared directory³. It is symbolized by a t character in the execute position of others. Option C is incorrect because other users cannot create files in /comptia, as they have no permissions on the directory or any files within it. Option D is incorrect because root is not the only user who can create files in /comptia. Any user who has write permission on the directory can create files within it, such as members of the wheel group.

NEW QUESTION 96

A Linux engineer set up two local DNS servers (10.10.10.10 and 10.10.10.20) and was testing email connectivity to the local mail server using the mail command on a local machine when the following error appeared:

```
Send-mail: Cannot open mail:25
```

The local machine DNS settings are:

```
$ cat /etc/resolv.conf
nameserver 10.10.10.10 #web records
nameserver 10.10.10.20 #email records

Mail server: mail.example.com
```

Which of the following commands could the engineer use to query the DNS server to get mail server information?

- A. dig @example.com 10.10.10.20 a
- B. dig @10.10.10.20 example.com mx
- C. dig @example.com 10.10.10.20 ptr
- D. dig @10.10.10.20 example.com ns

Answer: B

Explanation:

The command `dig @10.10.10.20 example.com mx` will query the DNS server to get mail server information. The dig command is a tool for querying DNS servers and displaying the results. The @ option specifies the DNS server to query, in this case 10.10.10.20. The mx option specifies the type of record to query, in this case mail exchange (MX) records, which identify the mail servers for a domain. The domain name to query is example.com. This command will show the MX records for example.com from the DNS server 10.10.10.20. This is the correct command to use to accomplish the task. The other options are incorrect because

they either use the wrong syntax (@example.com 10.10.10.20 instead of @10.10.10.20 example.com), the wrong type of record (a or ptr instead of mx), or the wrong domain name (example.com ns instead of example.com mx). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 415.

NEW QUESTION 98

Users have been unable to save documents to /home/tmp/temp and have been receiving the following error:

Path not found

A junior technician checks the locations and sees that /home/tmp/tempa was accidentally created instead of /home/tmp/temp. Which of the following commands should the technician use to fix this issue?

- A. cp /home/tmp/tempa /home/tmp/temp
- B. mv /home/tmp/tempa /home/tmp/temp
- C. cd /temp/tmp/tempa
- D. ls /home/tmp/tempa

Answer: B

Explanation:

The mv /home/tmp/tempa /home/tmp/temp command will fix the issue of the misnamed directory. This command will rename the directory /home/tmp/tempa to /home/tmp/temp, which is the expected path for users to save their documents. The cp /home/tmp/tempa /home/tmp/temp command will not fix the issue, as it will copy the contents of /home/tmp/tempa to a new file named /home/tmp/temp, not a directory. The cd /temp/tmp/tempa command will not fix the issue, as it will change the current working directory to /temp/tmp/tempa, which does not exist. The ls /home/tmp/tempa command will not fix the issue, as it will list the contents of /home/tmp/tempa, not rename it. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Files and Directories, page 413.

NEW QUESTION 99

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 104

Which of the following tools is commonly used for creating CI/CD pipelines?

- A. Chef
- B. Puppet
- C. Jenkins
- D. Ansible

Answer: C

Explanation:

The tool that is commonly used for creating CI/CD pipelines is Jenkins. Jenkins is an open-source automation server that enables continuous integration and continuous delivery (CI/CD) of software projects. Jenkins allows developers to build, test, and deploy code changes automatically and frequently using various plugins and integrations. Jenkins also supports distributed builds, parallel execution, pipelines as code, and real-time feedback. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Source Code; Jenkins

NEW QUESTION 107

A newly created container has been unable to start properly, and a Linux administrator is analyzing the cause of the failure. Which of the following will allow the administrator to determine the FIRST command that is executed inside the container right after it starts?

- A. docker export <container_id>
- B. docker info <container_id>
- C. docker start <container_id>
- D. docker inspect <container_id>

Answer: D

Explanation:

The command that will allow the administrator to determine the first command that is executed inside the container right after it starts is docker inspect <container_id>. This command will display detailed information about the container, including its configuration, state, network settings, mounts, and logs. One of the configuration fields is "Entrypoint", which shows the command that is executed when the container is run. The entrypoint can be specified in the Dockerfile or overridden at runtime using the --entrypoint option.

The other options are not correct commands for determining the first command that is executed inside the container. The docker export <container_id> command will export the contents of the container's filesystem as a tar archive to STDOUT. This will not show the entrypoint of the container, but only its files. The docker info <container_id> command is invalid because docker info does not take any arguments. It shows system-wide information about Docker, such as the number of containers, images, volumes, networks, and storage drivers. The docker start <container_id> command will start a stopped container and attach its STDOUT and

STDERR to the terminal. This will not show the endpoint of the container, but only its output. References: docker inspect | Docker Docs; docker export | Docker Docs; docker info | Docker Docs; docker start | Docker Docs

NEW QUESTION 112

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_logindf.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_logindf.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 116

Following the migration from a disaster recovery site, a systems administrator wants a server to require a user to change credentials at initial login. Which of the following commands should be used to ensure the aging attribute?

- A. chage -d 2 user
- B. chage -d 0 user
- C. chage -E 0 user
- D. chage -d 1 user

Answer: B

Explanation:

The chage command can be used to change the user password expiry information. The -d or --lastday option sets the last password change date. If the value is 0, the user will be forced to change the password at the next login. See chage command in Linux with examples and 10 chage command examples in Linux.

NEW QUESTION 119

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 && systemctl 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 && systemctl 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 && systemctl restart systemd-journald.service or cat /etc/systemd/journald.conf | awk '(print \$1,\$3)') or do not enable the Storage option (journalctl --list-boots && systemctl restart systemd-journald.service). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Managing Logging and Monitoring, page 489.

NEW QUESTION 121

An administrator is trying to diagnose a performance issue and is reviewing the following output:

```
avg-cpu:  %user  %nice  %system  %iowait  %steal   %idle
           2.00   0.00   3.00    32.00    0.00   63.00

Device            tps    kB_read/s    kB_wrtn/s    kB_read    kB_wrtn
sdb                345.00         0.02         0.04 4739073123 23849523
sdb1               345.00    32102.03    12203.01 4739073123 23849523
```

System Properties: CPU: 4 vCPU
 Memory: 40GB

Disk maximum IOPS: 690
 Disk maximum throughput: 44Mbps | 44000Kbps
 Based on the above output, which of the following BEST describes the root cause?

- A. The system has reached its maximum IOPS, causing the system to be slow.
- B. The system has reached its maximum permitted throughput, therefore iowait is increasing.
- C. The system is mostly idle, therefore the iowait is high.
- D. The system has a partitioned disk, which causes the IOPS to be doubled.

Answer: B

Explanation:

The system has reached its maximum permitted throughput, therefore iowait is increasing. The output of `iostat -x` shows that the device `sda` has an average throughput of 44.01 MB/s, which is equal to the disk maximum throughput of 44 Mbps. The output also shows that the device `sda` has an average iowait of 99.99%, which means that the CPU is waiting for the disk to complete the I/O requests. This indicates that the disk is the bottleneck and the system is slow due to the high iowait. The other options are incorrect because they are not supported by the outputs. The system has not reached its maximum IOPS, as the device `sda` has an average IOPS of 563.50, which is lower than the disk maximum IOPS of 690. The system is not mostly idle, as the output of `top` shows that the CPU is 100% busy. The system does not have a partitioned disk, as the output of `lsblk` shows that the device `sda` has only one partition `sda1`. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Optimizing Linux Systems, pages 513-514.

NEW QUESTION 126

A Linux systems administrator is setting up a new web server and getting 404 - NOT FOUND errors while trying to access the web server pages from the browser. While working on the diagnosis of this issue, the Linux systems administrator executes the following commands:

```
# getenforce
Enforcing

# matchpathcon -V /var/www/html/*
/var/www/html/index.html has context unconfined_u:object_r:user_home_t:s0, should be system_u:object_r:httpd_sys_content_t:s0
/var/www/html/page1.html has context unconfined_u:object_r:user_home_t:s0, should be system_u:object_r:httpd_sys_content_t:s0
```

Which of the following commands will BEST resolve this issue?

- A. `sed -i 's/SELINUX=enforcing/SELINUX=disabled/' /etc/selinux/config`
- B. `restorecon -R -v /var/www/html`
- C. `setenforce 0`
- D. `setsebool -P httpd_can_network_connect_db on`

Answer: B

Explanation:

The command `restorecon -R -v /var/www/html` will best resolve the issue. The issue is caused by the incorrect SELinux context of the web server files under the `/var/www/html` directory. The output of `ls -Z /var/www/html` shows that the files have the type `user_home_t`, which is not allowed for web content. The command `restorecon` restores the default SELinux context of files based on the policy rules. The options `-R` and `-v` are used to apply the command recursively and verbosely. This command will change the type of the files to `httpd_sys_content_t`, which is the correct type for web content. This will allow the web server to access the files and serve the pages to the browser. The other options are incorrect because they either disable SELinux entirely (`sed -i 's/SELINUX=enforcing/SELINUX=disabled/' /etc/selinux/config` or `setenforce 0`), which is not a good security practice, or enable an unnecessary boolean (`setsebool -P httpd_can_network_connect_db on`), which is not related to the issue. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 535.

NEW QUESTION 127

A systems administrator is tasked with mounting a USB drive on a system. The USB drive has a single partition, and it has been mapped by the system to the device `/dev/sdb`. Which of the following commands will mount the USB to `/media/usb`?

- A. `mount /dev/sdb1 /media/usb`
- B. `mount /dev/sdb0 /media/usb`
- C. `mount /dev/sdb /media/usb`
- D. `mount -t usb /dev/sdb1 /media/usb`

Answer: A

Explanation:

The `mount /dev/sdb1 /media/usb` command will mount the USB drive to `/media/usb`. This command will attach the filesystem on the first partition of the USB drive (`/dev/sdb1`) to the mount point `/media/usb`, making it accessible to the system. The `mount /dev/sdb0 /media/usb` command is invalid, as there is no such device as `/dev/sdb0`. The `mount /dev/sdb /media/usb` command is incorrect, as it will try to mount the entire USB drive instead of its partition, which may cause errors or data loss. The `mount -t usb /dev/sdb1 /media/usb` command is incorrect, as `usb` is not a valid filesystem type for mount. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Disk Storage, page 455.

NEW QUESTION 130

A Linux administrator needs to transfer a local file named `accounts.pdf` to a remote `/tmp` directory of a server with the IP address `10.10.10.80`. Which of the following commands needs to be executed to transfer this file?

- A. `rsync user@10.10.10.80: /tmp accounts.pdf`
- B. `scp accounts.pdf user@10.10.10.80:/tmp`
- C. `cp user@10.10.10.80: /tmp accounts.pdf`
- D. `ssh accounts.pdf user@10.10.10.80: /tmp`

Answer: B

Explanation:

The best command to use to transfer the local file accounts.pdf to the remote /tmp directory of the server with the IP address 10.10.10.80 is B. scp accounts.pdf user@10.10.10.80:/tmp. This command will use the secure copy protocol (scp) to copy the file from the local machine to the remote server over SSH. The command requires the username and password of the user on the remote server, as well as the full path of the destination directory.

The other commands are either incorrect or not suitable for this task. For example:

? A. rsync user@10.10.10.80:/tmp accounts.pdf will try to use the rsync command to synchronize files between the local and remote machines, but it has the wrong syntax and order of arguments. The source should come before the destination, and a colon (:) should separate the remote host and path.

? C. cp user@10.10.10.80:/tmp accounts.pdf will try to use the cp command to copy files, but it does not work over SSH and it has the wrong syntax and order of arguments. The source should come before the destination, and a colon (:) should separate the remote host and path.

? D. ssh accounts.pdf user@10.10.10.80:/tmp will try to use the ssh command to log into the remote server, but it has the wrong syntax and arguments. The username should come before the remote host, and a file name is not a valid argument for ssh.

NEW QUESTION 132

A Linux administrator needs to determine whether a hostname is in the DNS. Which of the following would supply the information that is needed?

- A. nslookup
- B. rsyn
- C. netstat
- D. host

Answer: A

Explanation:

The commands nslookup or host can be used to determine whether a hostname is in the DNS. The DNS is the domain name system, which is a service that translates domain names into IP addresses and vice versa. The nslookup command is a tool for querying the DNS and obtaining information about a domain name or an IP address. The host command is a similar tool that performs DNS lookups. Both commands can be used to check if a hostname is in the DNS by providing the hostname as an argument and seeing if the command returns a valid IP address or an error message. For example, the command nslookup www.google.com or host www.google.com will return the IP address of the Google website, while the command nslookup www.nosuchdomain.com or host www.nosuchdomain.com will return an error message indicating that the hostname does not exist. These commands will supply the information that is needed to determine whether a hostname is in the DNS. These are the correct commands to use for this task. The other options are incorrect because they do not query the DNS or obtain information about a hostname (rsync or netstat). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 378.

NEW QUESTION 133

What is the main objective when using Application Control?

- A. To filter out specific content.
- B. To assist the firewall blade with handling traffic.
- C. To see what users are doing.
- D. Ensure security and privacy of information.

Answer: D

Explanation:

The main objective when using Application Control is to ensure the security and privacy of information. Application Control is a security practice that blocks or restricts unauthorized applications from executing in ways that put data at risk. The control functions vary based on the business purpose of the specific application, but the main objective is to help ensure the privacy and security of data used by and transmitted between applications¹. Application Control can also prevent malware, untrusted, or unwanted applications from running on the network, reducing the risks and costs associated with data breaches¹. Application Control can also improve the overall network stability and performance by eliminating unnecessary or harmful applications¹. Application Control is not mainly used to filter out specific content, although it can be combined with other technologies such as URL filtering or content filtering to achieve that goal. Application Control is not mainly used to assist the firewall blade with handling traffic, although it can be integrated with firewall policies to enforce granular access rules based on applications. Application Control is not mainly used to see what users are doing, although it can provide visibility and reporting on application usage and activity.

NEW QUESTION 138

An administrator attempts to connect to a remote server by running the following command:

```
$ nmap 192.168.10.36
```

```
Starting Nmap 7.60 ( https://nmap.org ) at 2022-03-29 20:20 UTC Nmap scan report for www1 (192.168.10.36)
```

```
Host is up (0.000091s latency). Not shown: 979 closed ports PORT STATE SERVICE 21/tcp open ftp 22/tcp filtered ssh 631/tcp open ipp
```

```
Nmap done: 1 IP address (1 host up) scanned in 0.06 seconds
```

Which of the following can be said about the remote server?

- A. A firewall is blocking access to the SSH server.
- B. The SSH server is not running on the remote server.
- C. The remote SSH server is using SSH protocol version 1.
- D. The SSH host key on the remote server has expired.

Answer: A

Explanation:

This is because the port 22/tcp is shown as filtered by nmap, which means that nmap cannot determine whether the port is open or closed because a firewall or other device is blocking its probes. If the SSH server was not running on the remote server, the port would be shown as closed, which means that nmap received a TCP RST packet in response to its probe. If the remote SSH server was using SSH protocol version 1, the port would be shown as open, which means that nmap received a TCP SYN/ACK packet in response to its probe. If the SSH host key on the remote server had expired, the port would also be shown as open, but the SSH client would display a warning message about the host key verification failure. Therefore, the best explanation for the filtered state of the port 22/tcp is that a firewall is preventing nmap from reaching the SSH server.

You can find more information about nmap port states and how to interpret them in the following web search results:

? Nmap scan what does STATE=filtered mean?

? How to find ports marked as filtered by nmap

? Technical Tip: NMAP scan shows ports as filtered

NEW QUESTION 140

A systems administrator notices the process list on a mission-critical server has a large number of processes that are in state "Z" and marked as "defunct." Which of the following should the administrator do in an attempt to safely remove these entries from the process list?

- A. Kill the process with PID 1.
- B. Kill the PID of the processes.
- C. Kill the parent PID of the processes.
- D. Reboot the server.

Answer: C

Explanation:

As the web search results show, processes in state Z are defunct or zombie processes, which means they have terminated but their parent process has not reaped them properly. They do not consume any resources, but they occupy a slot in the process table. To remove them from the process list, the administrator needs to kill the parent process of the zombies, which will cause them to be reaped by the init process (PID 1). Killing the zombies themselves or the init process will not have any effect, as they are already dead. Rebooting the server may work, but it is not a safe or efficient option, as it may cause unnecessary downtime or data loss for a mission-critical server.

References

- ? Processes in a Zombie (Z) or Defunct State | Support | SUSE, paragraph 3
- ? linux - Zombie vs Defunct processes? - Stack Overflow, answer by admirableadmin
- ? How To Kill Zombie Processes on Linux | Linux Journal, paragraph 4

NEW QUESTION 141

An administrator deployed a Linux server that is running a web application on port 6379/tcp. SELinux is in enforcing mode based on organization policies. The port is open on the firewall. Users who are trying to connect to a local instance of the web application receive Error 13, Permission denied. The administrator ran some commands that resulted in the following output:

```
# semanage port -l | egrep '^http_port_t|6379'
http_port_t tcp 80, 81, 443, 488, 8008, 8009, 8443, 9000

# curl http://localhost/App.php
Cannot connect to App Server.
```

Which of the following commands should be used to resolve the issue?

- A. `semanage port -d -t http_port_t -p tcp 6379`
- B. `semanage port -a -t http_port_t -p tcp 6379`
- C. `semanage port -a http_port_t -p top 6379`
- D. `semanage port -l -t http_port_tcp 6379`

Answer: B

Explanation:

The command `semanage port -a -t http_port_t -p tcp 6379` adds a new port definition to the SELinux policy and assigns the type `http_port_t` to the port 6379/tcp. This allows the web application to run on this port and accept connections from users. This is the correct way to resolve the issue. The other options are incorrect because they either delete a port definition (-d), use the wrong protocol (top instead of tcp), or list the existing port definitions (-l). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 535.

NEW QUESTION 146

The development team created a new branch with code changes that a Linux administrator needs to pull from the remote repository. When the administrator looks for the branch in Git, the branch in question is not visible. Which of the following commands should the Linux administrator run to refresh the branch information?

- A. `git fetch`
- B. `git checkout`
- C. `git clone`
- D. `git branch`

Answer: A

Explanation:

The `git fetch` command downloads commits, files, and refs from a remote repository into the local one. It also updates the remote-tracking branches, which are references to the state of the remote branches. By running `git fetch`, the administrator can see the new branch created by the development team and then use `git checkout` to switch to it. References: 1: Git - git-fetch Documentation 2: Git Fetch | Atlassian Git Tutorial

NEW QUESTION 148

A Linux administrator has set up a new DNS forwarder and is configuring all internal servers to use the new forwarder to look up external DNS requests. The administrator needs to modify the firewall on the server for the DNS forwarder to allow the internal servers to communicate to it and make the changes persistent between server reboots. Which of the following commands should be run on the DNS forwarder server to accomplish this task?

- A. `ufw allow out dns`
- B. `systemctl reload firewalld`
- C. `iptables -A OUTPUT -p udp -ra udp -dport 53 -j ACCEPT`
- D. `firewall-cmd --zone-public --add-port-53/udp --permanent`

Answer: D

Explanation:

The command that should be run on the DNS forwarder server to accomplish the task is `firewall-cmd --zone=public --add-port=53/udp --permanent`.

The `firewall-cmd` command is a tool for managing `firewalld`, which is a firewall service that provides dynamic and persistent network security on Linux systems. The `firewalld` uses zones and services to define the rules and policies for the network traffic. The zones are logical groups of network interfaces and sources that have the same level of trust and security. The services are predefined sets of ports and protocols that are associated with certain applications or functions. The `--zone=public` option specifies the zone name that the rule applies to. The public zone is the default zone that represents the untrusted network, such as the internet. The `--add-port=53/udp` option adds a port and protocol to the zone. The 53 is the port number that is used by the DNS service. The `udp` is the protocol that is used by the DNS service. The `--permanent` option makes the change persistent across reboots. The command `firewall-cmd --zone=public --add-port=53/udp --permanent` will modify the firewall on the server for the DNS forwarder to allow the internal servers to communicate to it and make the changes persistent between server reboots. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not modify the firewall on the server for the DNS forwarder (`ufw allow out dns` or `systemctl reload firewalld`) or do not use the correct syntax for the command (`iptables -A OUTPUT -p udp -ra udp -dport 53 -j ACCEPT` instead of `iptables -A OUTPUT -p udp -ra udp --dport 53 -j ACCEPT`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 392.

NEW QUESTION 153

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`

References:

? [CompTIA Linux+ Study Guide], Chapter 11: Working with Containers, Section: Managing Containers with Docker

? [Docker PS Command with Examples]

NEW QUESTION 155

A Linux system is failing to boot with the following error:

```
error: no such partitions
Entering rescue mode...
grub rescue>
```

Which of the following actions will resolve this issue? (Choose two.)

- A. Execute `grub-install --root-directory=/mnt` and reboot.
- B. Execute `grub-install /dev/sdX` and reboot.
- C. Interrupt the boot process in the GRUB menu and add `rescue` to the kernel line.
- D. Fix the partition modifying `/etc/default/grub` and reboot.
- E. Interrupt the boot process in the GRUB menu and add `single` to the kernel line.
- F. Boot the system on a LiveCD/ISO.

Answer: BF

Explanation:

The administrator should do the following two actions to resolve the issue:

? Boot the system on a LiveCD/ISO. This is necessary to access the system and repair the boot loader. A LiveCD/ISO is a bootable media that contains a Linux distribution that can run without installation. The administrator can boot the system from the LiveCD/ISO and mount the root partition of the system to a temporary directory, such as `/mnt`.

? Execute `grub-install /dev/sdX` and reboot. This will reinstall the GRUB boot loader to the disk device, where `sdX` is the device name of the disk, such as `sda` or `sdb`. The GRUB boot loader is a program that runs when the system is powered on and allows the user to choose which operating system or kernel to boot. The issue is caused by a corrupted or missing GRUB boot loader, which prevents the system from booting. The command `grub-install` will restore the GRUB boot loader and fix the issue.

The other options are incorrect because they either do not fix the boot loader (interrupt the boot process in the GRUB menu or fix the partition modifying `/etc/default/grub`) or do not use the correct syntax (`grub-install --root-directory=/mnt` instead of `grub-install /dev/sdX` or `rescue` or `single` instead of `recovery` in the GRUB

menu). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing the Linux Boot Process, pages 265-266.

NEW QUESTION 158

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 162

A systems administrator creates a public key for authentication. Which of the following tools is most suitable to use when uploading the key to the remote servers?

- A. scp
- B. ssh-copy-id
- C. ssh-agent
- D. ssh-keyscan

Answer: B

Explanation:

The best tool to use when uploading the public key to the remote servers is

* B. `ssh-copy-id`. This tool will copy the public key from the local computer to the remote server and append it to the `authorized_keys` file, which is used for public key authentication. This tool will also create the necessary directories and files on the remote server if they do not exist. The other tools are either not suitable or not relevant for this task. For example:

? A. `scp` is a tool for securely copying files between hosts, but it does not automatically add the public key to the `authorized_keys` file.

? C. `ssh-agent` is a tool for managing private keys and passphrases, but it does not upload the public key to the remote server.

? D. `ssh-keyscan` is a tool for collecting public keys from remote hosts, but it does not upload the public key to the remote server.

NEW QUESTION 165

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

Output 1:

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

Output 2:

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

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

- A. Enable the `logsearch.service` and restart the service.
- B. Increase the `TimeoutStartUSec` configuration for the `logsearch.service`.
- C. Update the `OnCalendar` configuration to schedule the start of the `logsearch.service`.
- D. Update the `KillSignal` configuration for the `logsearch.service` to use `TERM`.

Answer: B

Explanation:

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

NEW QUESTION 168

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 173

A Linux administrator reviews a set of log output files and needs to identify files that contain any occurrence of the word denied. All log files containing entries in uppercase or lowercase letters should be included in the list. Which of the following commands should the administrator use to accomplish this task?

- A. `find . -type f -print | xargs grep -ln denied`
- B. `find . -type f -print | xargs grep -nv denied`
- C. `find . -type f -print | xargs grep -wL denied`
- D. `find . -type f -print | xargs grep -li denied`

Answer: D

Explanation:

The command `find . -type f -print | xargs grep -li denied` will accomplish the task of identifying files that contain any occurrence of the word denied. The `find` command is a tool for searching for files and directories on Linux systems. The `.` is the starting point of the search, which means the current directory. The `-type f` option specifies the type of the file, which means regular file. The `-print` option prints the full file name on the standard output. The `|` is a pipe symbol that redirects the output of one command to the input of another command. The `xargs` command is a tool for building and executing commands from standard input. The `grep` command is a tool for searching for patterns in files or input.

The `-li` option specifies the flags that the `grep` command should apply. The `-l` flag shows only the file names that match the pattern, instead of the matching lines. The `-i` flag ignores the case of the pattern, which means it matches both uppercase and lowercase letters.

The `denied` is the pattern that the `grep` command should search for. The command `find . -type f -print | xargs grep -li denied` will find all the regular files in the current directory and its subdirectories, and then search for any occurrence of the word denied in those files, ignoring the case, and print only the file names that match the pattern. This will allow the administrator to identify files that contain any occurrence of the word denied. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not ignore the case of the pattern (`find . -type f -print | xargs grep -ln denied` or `find . -type f -print | xargs grep -wL denied`) or do not show the file names that match the pattern (`find . -type f -print | xargs grep -nv denied`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Managing Logging and Monitoring, page 489.

NEW QUESTION 178

An administrator would like to list all current containers, regardless of their running state. Which of the following commands would allow the administrator to accomplish this task?

- A. `docker ps -a`
- B. `docker list`
- C. `docker image ls`
- D. `docker inspect image`

Answer: A

Explanation:

The best command to use to list all current containers, regardless of their running state, is `A. docker ps -a`. This command will show all containers, both running and stopped, with details such as container ID, image name, status, and ports. The other commands are either invalid or not relevant for this task. For example:

? `B. docker list` is not a valid command. There is no subcommand named `list` in `docker`.

? `C. docker image ls` will list all the images available on the local system, not the containers.

? `D. docker inspect image` will show detailed information about a specific image, not all the containers.

NEW QUESTION 182

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 185

A Linux system is having issues. Given the following outputs:

```
# dig @192.168.2.2 mycompthost
```

```
;<< >> DiG 9.9.4-RedHat-9.9.4-74.el7_6.1 << >> @192.168.2.2 mycomptiahost
;(1 server found)
;; global options: +cmd
;; connection timed out; no servers could be reached
# nc -v 192.168.2.2 53
Ncat: Version 7.70 ( https://nmap.org/ncat ) Ncat: Connection timed out.
# ping 192.168.2.2
PING 192.168.2.2 (192.168.2.2) 56(84) bytes of data.
64 bytes from 192.168.2.2: icmp_seq=1 ttl=117 time=4.94 ms 64 bytes from 192.168.2.2: icmp_seq=2 ttl=117 time=10.5 ms Which of the following best describes this issue?
```

- A. The DNS host is down.
- B. The name mycomptiahost does not exist in the DNS.
- C. The Linux engineer is using the wrong DNS port.
- D. The DNS service is currently not available or the corresponding port is blocked.

Answer: D

Explanation:

The ping command shows that the Linux system can reach the DNS server at 192.168.2.2, so the DNS host is not down. The dig and nc commands show that the Linux system cannot connect to the DNS server on port 53, which is the standard port for DNS queries. This means that either the DNS service is not running on the DNS server, or there is a firewall or network device blocking the port 53 traffic. Therefore, the DNS service is currently not available or the corresponding port is blocked. References: 1: How To Troubleshoot DNS Client Issues in Linux - RootUsers 2: 6 Best Tools to Troubleshoot DNS Issues in Linux - Tecmint 3: How To Troubleshoot DNS in Linux - OrcaCore 4: Fixing DNS Issues in Ubuntu 20.04 | DeviceTests

NEW QUESTION 190

A systems administrator is deploying three identical, cloud-based servers. The administrator is using the following code to complete the task:

```
resource "aws_instance" "ec2_instance" {
  ami                = data.aws_ami.vendor-Linux-2.id
  associate_public_ip_address = true
  count              = 3
  instance_type      = "instance_type"
  vpc_security_group_ids = [aws_security_group.allow_ssh.id]
  key_name            = aws_key_pair.key_pair.key_name

  tags = {
    Name = "${var.namespace} ${count.index}"
  }
}
```

Which of the following technologies is the administrator using?

- A. Ansible
- B. Puppet
- C. Chef
- D. Terraform

Answer: D

Explanation:

The code snippet is written in Terraform language, which is a tool for building, changing, and versioning infrastructure as code. Terraform uses a declarative syntax to describe the desired state of the infrastructure and applies the changes accordingly. The code defines a resource of type aws_instance, which creates an AWS EC2 instance, and sets the attributes such as the AMI ID, instance type, security group IDs, and key name. The code also uses a count parameter to create three identical instances and assigns them different names using the count.index variable. This is the correct technology that the administrator is using. The other options are incorrect because they use different languages and syntaxes for infrastructure as code. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 559.

NEW QUESTION 192

A systems administrator is trying to track down a rogue process that has a TCP listener on a network interface for remote command-and-control instructions. Which of the following commands should the systems administrator use to generate a list of rogue process names? (Select two).

- A. netstat -antp | grep LISTEN
- B. lsof -iTCP | grep LISTEN
- C. lsof -i:22 | grep TCP
- D. netstat -a | grep TCP
- E. nmap -p1-65535 | grep -i tcp
- F. nmap -sS 0.0.0.0/0

Answer: AB

Explanation:

The best commands to use to generate a list of rogue process names that have a TCP listener on a network interface are A. netstat -antp | grep LISTEN and B. lsof -iTCP | grep LISTEN. These commands will show the process ID (PID) and name of the processes that are listening on TCP ports, which can be used to

identify any suspicious or unauthorized processes. The other commands are either not specific enough, not valid, or not relevant for this task. For example:
 ? C. `lsof -i:22 | grep TCP` will only show the processes that are listening on port 22, which is typically used for SSH, and not any other ports.
 ? D. `netstat -a | grep TCP` will show all the TCP connections, both active and listening, but not the process names or IDs.
 ? E. `nmap -p1-65535 | grep -i tcp` will scan all the TCP ports on the local host, but not show the process names or IDs.
 ? F. `nmap -sS 0.0.0.0/0` will perform a stealth scan on the entire internet, which is not only impractical, but also illegal in some countries.

NEW QUESTION 195

A Linux administrator is troubleshooting an issue in which users are not able to access `https://portal.comptia.org` from a specific workstation. The administrator runs a few commands and receives the following output:

```
# cat /etc/hosts
10.10.10.55 portal.comptia.org

# host portal.comptia.org
portal.comptia.org has address 192.168.1.55

#cat /etc/resolv.conf
nameserver 10.10.10.5
```

Which of the following tasks should the administrator perform to resolve this issue?

- A. Update the name server in `resolv.conf` to use an external DNS server.
- B. Remove the entry for `portal.comptia.org` from the local hosts file.
- C. Add a network route from the `10.10.10.0/24` to the `192.168.0.0/16`.
- D. Clear the local DNS cache on the workstation and rerun the `host` command.

Answer: B

Explanation:

The best task to perform to resolve this issue is B. Remove the entry for `portal.comptia.org` from the local hosts file. This is because the local hosts file has a wrong entry that maps `portal.comptia.org` to `10.10.10.55`, which is different from the actual IP address of `192.168.1.55` that is returned by the DNS server. This causes a mismatch and prevents the workstation from accessing the website. By removing or correcting the entry in the hosts file, the workstation will use the DNS server to resolve the domain name and access the website successfully.

To remove or edit the entry in the hosts file, you need to have root privileges and use a text editor such as `vi` or `nano`. For example, you can run the command:
`sudo vi /etc/hosts`
 and delete or modify the line that says: `10.10.10.55 portal.comptia.org`
 Then save and exit the file.

NEW QUESTION 198

Which of the following files holds the system configuration for journal when running `systemd`?

- A. `/etc/systemd/journald.conf`
- B. `/etc/systemd/systemd-journalctl.conf`
- C. `/usr/lib/systemd/journalctl.conf`
- D. `/etc/systemd/systemd-journald.conf`

Answer: A

Explanation:

The file that holds the system configuration for journal when running `systemd` is `/etc/systemd/journald.conf`. This file contains various settings that control the behavior of the `journald` daemon, which is responsible for collecting and storing log messages from various sources. The `journald.conf` file can be edited to change the default values of these settings, such as the storage location, size limits, compression, and forwarding options of the journal files. The file also supports a drop-in directory `/etc/systemd/journald.conf.d/` where additional configuration files can be placed to override or extend the main file. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks; `journald.conf(5)` - Linux manual page

NEW QUESTION 201

A junior systems administrator has just generated public and private authentication keys for passwordless login. Which of the following files will be moved to the remote servers?

- A. `id_dsa.pem`
- B. `id_rsa`
- C. `id_ecdsa`
- D. `id_rsa.pub`

Answer: D

Explanation:

The file `id_rsa.pub` will be moved to the remote servers for passwordless login. The `id_rsa.pub` file is the public authentication key that is generated by the `ssh-keygen` command. The public key can be copied to the remote servers by using the `ssh-copy-id` command or manually. The remote servers will use the public key to authenticate the user who has the corresponding private key (`id_rsa`). This will allow the user to log in without entering a password. The other options are incorrect because they are either private keys (`id_rsa`, `id_dsa.pem`, or `id_ecdsa`) or non-existent files (`id_dsa.pem` or `id_ecdsa`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 410.

NEW QUESTION 202

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 204

A Linux administrator needs to create a new cloud.cpio archive containing all the files from the current directory. Which of the following commands can help to accomplish this task?

- A. `ls | cpio -iv > cloud.cpio`
- B. `ls | cpio -iv < cloud.cpio`
- C. `ls | cpio -ov > cloud.cpio`
- D. `ls cpio -ov < cloud.cpio`

Answer: C

Explanation:

The command `ls | cpio -ov > cloud.cpio` can help to create a new cloud.cpio archive containing all the files from the current directory. The `ls` command lists the files in the current directory and outputs them to the standard output. The `|` operator pipes the output to the next command. The `cpio` command is a tool for creating and extracting compressed archives. The `-o` option creates a new archive and the `-v` option shows the verbose output. The `>` operator redirects the output to the cloud.cpio file. This command will create a new cloud.cpio archive with all the files from the current directory. The other options are incorrect because they either use the wrong options (`-i` instead of `-o`), the wrong arguments (`cloud.epio` instead of `cloud.cpio`), or the wrong syntax (`<` instead of `>` or missing `|`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, page 351.

NEW QUESTION 209

Employees in the finance department are having trouble accessing the file /opt/work/file. All IT employees can read and write the file. Systems administrator reviews the following output:

```
admin@server:/opt/work$ ls -al file
-rw-rw----+ 1 root it 4 Sep 5 17:29 file
```

Which of the following commands would permanently fix the access issue while limiting access to IT and finance department employees?

- A. `chattr +i file`
- B. `chown it:finance file`
- C. `chmod 666 file`
- D. `setfacl -m g:finance:rw file`

Answer: D

Explanation:

The command `setfacl -m g:finance:rw file` will permanently fix the access issue while limiting access to IT and finance department employees. The `setfacl` command is a tool for modifying the access control lists (ACLs) of files and directories on Linux systems. The ACLs are a mechanism that allows more fine-grained control over the permissions of files and directories than the traditional owner-group-others model. The `-m` option specifies the modification to the ACL. The `g:finance:rw` means that the group named finance will have read and write permissions on the file. The file is the name of the file to modify, in this case /opt/work/file. The command `setfacl -m g:finance:rw file` will add an entry to the ACL of the file that will grant read and write access to the finance group. This will fix the access issue and allow the finance employees to access the file. The command will also preserve the existing permissions of the file, which means that the IT employees will still have read and write access to the file. This will limit the access to IT and finance department employees and prevent unauthorized access from other users.

This is the correct command to use to accomplish the task. The other options are incorrect because they either do not fix the access issue (`chattr +i file` or `chown it:finance file`) or do not limit the access to IT and finance department employees (`chmod 666 file`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing File Permissions and Ownership, page 352.

NEW QUESTION 211

Which of the following technologies provides load balancing, encryption, and observability in containerized environments?

- A. Virtual private network
- B. Sidecar pod
- C. Overlay network
- D. Service mesh

Answer: D

Explanation:

"A service mesh controls the delivery of service requests in an application. Common features provided by a service mesh include service discovery, load balancing, encryption and failure recovery."

The technology that provides load balancing, encryption, and observability in containerized environments is service mesh. A service mesh is a dedicated infrastructure layer that manages the communication and security between microservices in a distributed system. A service mesh consists of two components: a data plane and a control plane. The data plane is composed of proxies that are deployed alongside the microservices as sidecar pods. The proxies handle the network traffic between the microservices and provide features such as load balancing, encryption, authentication, authorization, routing, and observability. The control plane is responsible for configuring and managing the data plane and providing a unified interface for the administrators and developers. A service mesh can help improve the performance, reliability, and security of containerized applications and simplify the development and deployment process. A service mesh is the technology that provides load balancing, encryption, and observability in containerized environments. This is the correct answer to the question. The other options are incorrect because they either do not provide all the features of a service mesh (virtual private network or overlay network) or are not a technology but a component of a service mesh (sidecar pod). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 574. <https://www.techtarget.com/searchitoperations/definition/service-mesh>

NEW QUESTION 212

After listing the properties of a system account, a systems administrator wants to remove the expiration date of a user account. Which of the following commands will accomplish this task?

- A. `chgrp system accountname`
- B. `passwd -s accountname`
- C. `chmod -G system account name`
- D. `chage -E -1 accountname`

Answer: D

Explanation:

The command `chage -E -1 accountname` will accomplish the task of removing the expiration date of a user account. The `chage` command is a tool for changing user password aging information on Linux systems. The `-E` option sets the expiration date of the user account, and the `-1` value means that the account will never expire. The command `chage -E -1 accountname` will remove the expiration date of the user account named `accountname`. This is the correct command to use to accomplish the task. The

other options are incorrect because they either do not affect the expiration date (`chgrp`, `passwd`, or `chmod`) or do not exist (`chmod -G`). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Users and Groups, page 467.

NEW QUESTION 215

An administrator started a long-running process in the foreground that needs to continue without interruption. Which of the following keystrokes should the administrator use to continue running the process in the background?

- A. `<Ctrl+z> bg`
- B. `<Ctrl+d> bg`
- C. `<Ctrl+b> jobs -1`
- D. `<Ctrl+h> bg &`

Answer: A

Explanation:

A long-running process is a program that takes a long time to complete or runs indefinitely on a Linux system. A foreground process is a process that runs in the current terminal and receives input from the keyboard and output to the screen. A background process is a process that runs in the background and does not interact with the terminal. A background process can continue running even if the terminal is closed or disconnected.

To start a long-running process in the background, the user can append an ampersand (`&`) to the command, such as `someapp &`. This will run `someapp` in the background and return control to the terminal immediately.

To move a long-running process from the foreground to the background, the user can use two keystrokes: `Ctrl+Z` and `bg`. The `Ctrl+Z` keystroke will suspend (pause) the foreground process and return control to the terminal. The `bg` keystroke will resume (continue) the suspended process in the background and detach it from the terminal. The statement B is correct.

The statements A, C, and D are incorrect because they do not perform the desired task. The `bg` keystroke alone will not work unless there is a suspended process to resume. The `Ctrl+B` keystroke will not suspend the foreground process, but rather move one character backward in some applications. The `jobs` keystroke will list all processes associated with the current terminal. The `bg &` keystroke will cause an error because `bg` does not take any arguments. References: [How to Run Linux Processes in Background]

NEW QUESTION 219

An administrator added the port 2222 for the SSH server on myhost and restarted the SSH server. The administrator noticed issues during the startup of the service. Given the following outputs:

```
$ ssh -p 2222 myhost
ssh:connect to host myhost on port 2222: Connection refused

$ nmap -p 2222 myhost
Starting Nmap 7.70 ( https://nmap.org ) at 2022-10-17 21:12 EEST
Nmap scan report for myhost (10.7.3.26)
Host is up (0.00027s latency).
rDNS record for 10.7.3.26: myhost
PORT      STATE SERVICE
2222/tcp  closed EtherNetIP-1
MAC Address: 52:54:00:F5:DF:F8 (QEMU virtual NIC)
Nmap done: 1 IP address (1 host up) scanned in 0.57 seconds

$ systemctl status sshd
   • sshd.service - OpenSSH server daemon
  Loaded: loaded (/usr/lib/systemd/system/sshd.service; enabled; vendor preset: enabled)
  Active: active (running) since Mon 2022-10-17 19:40:07 CEST; 36min ago
    Docs: man:sshd(8)
          man:sshd_config(5)
 Main PID: 13186 (sshd)
   Tasks: 1 (limit: 12373)
  Memory: 1.1M
  CGroup: /system.slice/sshd.service
          └─13186 /usr/sbin/sshd -D -oCiphers=aes256-gcm@openssh.com

Oct 17 19:40:07 myhost systemd[1]: Starting OpenSSH server daemon...
Oct 17 19:40:07 myhost sshd[13186]: error: Bind to port 2222 on 0.0.0.0 failed: Permission denied.
Oct 17 19:40:07 myhost systemd[1]: Started OpenSSH server daemon.
Oct 17 19:40:07 myhost sshd[13186]: Server listening on 0.0.0.0 port 22.
```

Which of the following commands will fix the issue?

- A. `semanage port -a -t ssh_port_t -p tcp 2222`
- B. `chcon system_u:object_r:ssh_home_t /etc/ssh/*`
- C. `iptables -A INPUT -p tcp -- dport 2222 -j ACCEPT`
- D. `firewall-cmd -- zone=public -- add-port=2222/tcp`

Answer: A

Explanation:

The correct answer is A. `semanage port -a -t ssh_port_t -p tcp 2222`

This command will allow the SSH server to bind to port 2222 by adding it to the SELinux policy. The `semanage` command is a utility for managing SELinux policies. The `port` subcommand is used to manage network port definitions. The `-a` option is used to add a new record, the `-t` option is used to specify the SELinux type, the `-p` option is used to specify the protocol, and the `tcp 2222` argument is used to specify the port number. The `ssh_port_t` type is the default type for SSH ports in SELinux.

The other options are incorrect because:

* B. `chcon system_u:object_r:ssh_home_t /etc/ssh/*`

This command will change the SELinux context of all files under `/etc/ssh/` to `system_u:object_r:ssh_home_t`, which is not correct. The `ssh_home_t` type is used for user home directories that are accessed by SSH, not for SSH configuration files. The correct type for SSH configuration files is `sshd_config_t`.

* C. `iptables -A INPUT -p tcp --dport 2222 -j ACCEPT`

This command will add a rule to the iptables firewall to accept incoming TCP connections on port 2222. However, this is not enough to fix the issue, as SELinux will still block the SSH server from binding to that port. Moreover, iptables may not be the default firewall service on some Linux distributions, such as Fedora or CentOS, which use `firewalld` instead.

* D. `firewall-cmd --zone=public --add-port=2222/tcp`

This command will add a rule to the `firewalld` firewall to allow incoming TCP connections on port 2222 in the public zone. However, this is not enough to fix the issue, as SELinux will still block the SSH server from binding to that port. Moreover, `firewalld` may not be installed or enabled on some Linux distributions, such as Ubuntu or Debian, which use iptables instead.

References:

- ? [How to configure SSH to use a non-standard port with SELinux set to enforcing](#)
- ? [Change SSH Port on CentOS/RHEL/Fedora With SELinux Enforcing](#)
- ? [How to change SSH port when SELinux policy is enabled](#)

NEW QUESTION 220

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions 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](#)