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 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 4
A Linux administrator is tasked with creating resources using containerization. When deciding how to create this type of deployment, the administrator identifies some key features, including portability, high availability, and scalability in production. Which of the following should the Linux administrator choose for the new design?

A. Docker
B. On-premises systems
C. Cloud-based systems
D. Kubernetes

Answer: D

Explanation:
The Linux administrator should choose Kubernetes for the new design that requires portability, high availability, and scalability in production using containerization. Kubernetes is an open-source platform that automates the deployment, scaling, and management of containerized applications across clusters of nodes. Kubernetes provides features such as service discovery, load balancing, storage orchestration, self-healing, secret and configuration management, and batch execution. Kubernetes also supports multiple container runtimes, such as Docker, containerd, and CRI-O, making it portable across different platforms and clouds. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks; What is Kubernetes? | Kubernetes

NEW QUESTION 5
In which of the following filesystems are system logs commonly stored?

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

Answer: A

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

NEW QUESTION 6
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 7
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 makemake 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 8
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 ap
B. conf.
C. Run git stash ap
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 9
A systems administrator wants to permit access temporarily to an application running on port 1234/TCP on a Linux server. Which of the following commands will permit this traffic?

A. firewall-cmd —new-service=1234/tcp
B. firewall-cmd —service=1234 —protocol=tcp
C. firewall-cmd —add—port=1234/tcp
D. firewall-cmd —add-whitelist-uid=1234

Answer: C

Explanation:
The firewall-cmd command is used to manage firewalld, which is a firewall service for Linux systems that provides dynamic and persistent configuration of firewall rules. Firewalld uses zones and services to define different levels of trust and access for network connections.
To permit access temporarily to an application running on port 1234/TCP on a Linux server, the systems administrator can use the firewall-cmd --add-port=1234/tcp command. This command will add a rule to the default zone (usually public) that allows incoming traffic on port 1234/TCP. The rule will only be effective until the next reload or restart of firewalld. To make the rule permanent, the administrator can add the --permanent option to the command. The statement C is correct.
The statements A, B, and D are incorrect because they do not permit access to port 1234/TCP. The firewall-cmd --new-service=1234/tcp command does not exist. The firewall- cmd --service=1234 --protocol=tcp command does not work because 1234 is not a predefined service name in firewalld. The firewall-cmd --add-whitelist-uid=1234 command does not exist. References: [How to Use FirewallD to Manage Firewall in Linux]

NEW QUESTION 10
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 10
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 hostnane kernel: iptables denied: IN=eth0 OUT=
MAC=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 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 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 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 14
A systems administrator needs to clone the partition /dev/sdc1 to /dev/sdd1. Which of the following commands will accomplish this task?

A. tar -cvzf /dev/sdd1 /dev/sdc1
B. rsync /dev/sdc1 /dev/sdd1
C. dd if=/dev/sdc1 of=/dev/sdd1
D. scp /dev/sdc1 /dev/sdd1

Answer: C

Explanation:
The command dd if=/dev/sdc1 of=/dev/sdd1 copies the data from the input file (if) /dev/sdc1 to the output file (of) /dev/sdd1, byte by byte. This is the correct way to clone a partition. The other options are incorrect because they either compress the data (tar -cvzf), synchronize the files (rsync), or copy the files over a network (scp), which are not the same as cloning a partition. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 321.

NEW QUESTION 16
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 20
A Linux administrator wants to prevent the httpd web service from being started both manually and automatically on a server. Which of the following should the administrator use to accomplish this task?

A. systemctl mask httpd
B. systemctl disable httpd
C. systemctl stop httpd
D. systemctl reload httpd

Answer: A

Explanation:
The best command to use to prevent the httpd web service from being started both manually and automatically on a server is A. systemctl mask httpd. This command will create a symbolic link from the httpd service unit file to /dev/null, which will make the service impossible to start or enable. This is different from systemctl disable httpd, which will only prevent the service from starting automatically on boot, but not manually. The other commands are either not relevant or not sufficient for this task. For example:
? C. systemctl stop httpd will only stop the service if it is currently running, but it will not prevent it from being started again.
? D. systemctl reload httpd will only reload the configuration files of the service, but it
will not stop or disable it.

NEW QUESTION 25
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 28
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 32
Users report that connections to a MariaDB service are being closed unexpectedly. A systems administrator troubleshoots the issue and finds the following message in /var/log/messages:

```
dbserver kernel: out of Memory: Killed process 1234 (mysqld).
```

Which of the following is causing the connection issue?

A. The process mysqld is using too many semaphores.
B. The server is running out of file descriptors.
C. Something is starving the server resources.
D. The amount of RAM allocated to the server is too high.

Answer: B

Explanation:
The message in /var/log/messages indicates that the server is running out of file descriptors. A file descriptor is a non-negative integer identifier for an open file in Linux. Each process has a table of open file descriptors where a new entry is appended upon opening a new file. There is a limit on how many file descriptors a process can open at a time, which depends on the system configuration and the user privileges. If a process tries to open more files than the limit, it will fail with an error message like "Too many open files". This could cause connections to be closed unexpectedly or other problems with the application.
The other options are not correct causes for the connection issue. The process mysqld is not using too many semaphores, which are synchronization mechanisms for processes that share resources. Semaphores are not related to file descriptors or open files. Something is not starving the server resources, which could mean high CPU usage, memory pressure, disk I/O, network congestion, or other factors that affect performance. These could cause slowdowns or timeouts, but not file descriptor exhaustion. The amount of RAM allocated to the server is not too high, which could cause swapping or paging if it exceeds the physical memory available. This could also affect performance, but not file descriptor availability. References: File Descriptor Requirements (Linux Systems); Limits on the Number of Linux File Descriptors

NEW QUESTION 33
Which of the following can be used as a secure way to access a remote termi-nal?

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 37
A junior administrator is trying to set up a passwordless SSH connection to one of the servers. The administrator follows the instructions and puts the key in the authorized_key file at the server, but the administrator is still asked to provide a password during the connection.
Given the following output:

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

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

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

Answer: B

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

NEW QUESTION 42
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 43
A developer needs to launch an Nginx image container, name it Web001, and ex-pose 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 Web00l 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 45
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 46
A Linux systems administrator is configuring a new filesystem that needs the capability to be mounted persistently across reboots. Which of the following commands will accomplish this task? (Choose two.)

A. df -h /data
B. mkfs.ext4 /dev/sdc1
C. fsck /dev/sdc1
D. fdisk -l /dev/sdc1
E. echo "/data /dev/sdc1 ext4 defaults 0 0" >> /etc/fstab
F. echo "/dev/sdc1 /data ext4 defaults 0 0" >> /etc/fstab

Answer: BF

Explanation:
"modify the /etc/fstab text file to automatically mount the new partition by opening it in an editor and adding the following line:
/dev/ xxx 1 /data ext4 defaults 1 2
where xxx is the device name of the storage device"
https://learning.oreilly.com/library/view/mastering-linux-system/9781119794455/b01.xhtml To configure a new filesystem that needs the capability to be mounted persistently across reboots, two commands are needed: mkfs.ext4 /dev/sdc1 and echo “/dev/sdc1 /data ext4 defaults 0 0” >> /etc/fstab. The first command creates an ext4 filesystem on the device /dev/sdc1, which is the partition that will be used for the new filesystem. The second command appends a line to the /etc/fstab file, which is the configuration file that controls persistent mount points of filesystems. The line specifies the device name, the mount point (/data), the filesystem type (ext4), the mount options (defaults), and the dump and pass values (0 0). The other commands are incorrect because they either do not create or configure a filesystem, or they have wrong syntax or arguments. References: CompTIA Linux+ Study Guide, Fourth Edition, page 409-410, 414-415.

NEW QUESTION 47
A Linux administrator needs to create an image named sda.img from the sda disk and store it in the /tmp directory. Which of the following commands should be used to accomplish this task?

A. dd of=/dev/sda if=/tmp/sda.img
B. dd if=/dev/sda of=/tmp/sda.img
C. dd --if=/dev/sda --of=/tmp/sda.img
D. dd --of=/dev/sda --if=/tmp/sda.img

Answer: B

Explanation:
The command dd if=/dev/sda of=/tmp/sda.img should be used to create an image named sda.img from the sda disk and store it in the /tmp directory. The dd command is a tool for copying and converting data on Linux systems. The if option specifies the input file or device, in this case /dev/sda, which is the disk device. The of option specifies the output file or device, in this case /tmp/sda.img, which is the image file. The command dd if=/dev/sda of=/tmp/sda.img will copy the entire

disk data from /dev/sda to /tmp/sda.img and create an image file. This is the correct command to use to accomplish the task. The other options are incorrect because they either use the wrong options (--if or --of instead of if or of) or swap the input and output (dd of=/dev/sda if=/tmp/sda.img or dd --of=/dev/sda --if=/tmp/sda.img). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 323.

NEW QUESTION 50
Users in the human resources department are trying to access files in a newly created directory. Which of the following commands will allow the users access to the files?

A. chattr
B. chgrp
C. chage
D. chcon

Answer: B

Explanation:
The chgrp command is used to change the group ownership of files and directories. By using this command, the administrator can assign the files in the newly created directory to the human resources group, which will allow the users in that group to access them. The other commands are not relevant for this task. For example:
? chattr is used to change the file attributes, such as making them immutable or append-only1.
? chage is used to change the password expiration information for a user account2.
? chcon is used to change the security context of files and directories, which is related to SELinux3.
References:
? The CompTIA Linux+ Certification Exam Objectives mention that the candidate should be able to “manage file and directory ownership and permissions” as part of the Hardware and System Configuration domain4.
? The web search result 2 explains how to use the chgrp command with examples.
? The web search result 3 compares the chmod and chgrp commands and their effects on file permissions.

NEW QUESTION 51
A Linux administrator copied a Git repository locally, created a feature branch, and committed some changes to the feature branch. Which of the following Git actions should the Linux administrator use to publish the changes to the main branch of the remote repository?

A. rebase
B. tag
C. commit
D. push

Answer: D

Explanation:
The push action is used to publish the changes made in a local branch to a remote branch of a Git repository. This action will update the remote branch with the commits made in the local branch and synchronize the two branches. The rebase action is used to reapply commits from one branch onto another branch, creating a linear history of commits. This action does not publish any changes to a remote repository. The tag action is used to create an annotated reference to a specific commit in a Git repository. This action does not publish any changes to a remote repository. The commit action is used to record changes made in the local repository and create a new snapshot of the project state. This action does not publish any changes to a remote repository. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 20: Writing and Executing Bash Shell Scripts, page 579.

NEW QUESTION 55
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 60
A systems administrator is tasked with preventing logins from accounts other than root, while the file /etc/nologin exists. Which of the following PAM modules will accomplish this task?

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

Answer: D

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

Managing Users and Groups, page 471.

NEW QUESTION 64
An administrator needs to make some changes in the IaC declaration templates. Which of the following commands would maintain version control?

A. git clone https://github.com/comptia/linux+-.git git push origin
B. git clone https://qithub.com/comptia/linux+-.git git fetch New-Branch
C. git clone https://github.com/comptia/linux+-.git git status
D. git clone https://github.com/comptia/linuxt+-.git git checkout -b <new-branch>

Answer: D

Explanation:
The command that will maintain version control while making some changes in the IaC declaration templates is git checkout -b <new-branch>. This command uses the git tool, which is a distributed version control system that tracks changes in source code and enables collaboration among developers. The checkout option switches to a different branch in the git repository, where a branch is a pointer to a specific commit in the history. The -b option creates a new branch with the given name, and switches to it. This way, the administrator can make changes in the new branch without affecting the main branch, and later merge them if needed.
The other options are not correct commands for maintaining version control while making some changes in the IaC declaration templates. The git clone https://github.com/comptia/linux±.git command will clone an existing repository from a remote URL to a local directory, but it will not create a new branch for making changes. The git push origin command will push the local changes to a remote repository named origin, but it will not create a new branch for making changes. The git fetch New-Branch command will fetch updates from a remote branch named New-Branch, but it will not create a new branch for making changes.
References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Source Code; Git - Basic Branching and Merging

NEW QUESTION 65
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 67
A Linux administrator found many containers in an exited state. Which of the following commands will allow the administrator to clean up the containers in an exited state?

A. docker rm -- all
B. docker rm $ (docker ps -aq)
C. docker images prune *
D. docker rm -- state exited

Answer: B

Explanation:
This command will remove all containers, regardless of their state, by passing the IDs of all containers to the docker rm command. The docker ps -aq command will list the IDs of all containers, including the ones in an exited state, and the $ ( ) syntax will substitute the output of the command as an argument for the docker rm command. This is a quick and easy way to clean up all containers, but it may also remove containers that are still needed or running.
References
? docker rm | Docker Docs - Docker Documentation, section “Remove all containers”
? Docker Remove Exited Containers | Easy methods. - Bobcares, section “For removing all exited containers”

NEW QUESTION 72
Using AD Query, the security gateway connections to the Active Directory Domain Controllers using what protocol?

A. Windows Management Instrumentation (WMI)
B. Hypertext Transfer Protocol Secure (HTTPS)
C. Lightweight Directory Access Protocol (LDAP)
D. Remote Desktop Protocol (RDP)

Answer: C

Explanation:
Using AD Query, the security gateway connects to the Active Directory Domain Controllers using Lightweight Directory Access Protocol (LDAP). LDAP is a protocol that provides access to directory services over a network. AD Query uses LDAP queries to retrieve information about users and groups from Active Directory Domain Controllers without installing any software on them. AD Query does not use Windows Management Instrumentation (WMI), Hypertext Transfer Protocol Secure (HTTPS), or Remote Desktop Protocol (RDP) to connect to Active Directory Domain Controllers. References: Check Point Certified Security Administrator (CCSA) R80.x Study Guide, Chapter 5: User Management and Authentication, page 69.

NEW QUESTION 75
A systems administrator checked out the code from the repository, created a new branch, made changes to the code, and then updated the main branch. The

systems administrator wants to ensure that the Terraform state files do not appear in the main branch. Which of following should the administrator use to meet this requirement?

A. clone
B. gitxgnore
C. get
D. .ssh

Answer: B

Explanation:
To prevent certain files from being tracked by Git, the administrator can use a .gitignore file (B) in the repository. The .gitignore file can specify patterns of files or directories that Git should ignore. This way, the Terraform state files will not appear in the main branch or any other branch. The other commands are not related to this requirement. References:
? [CompTIA Linux+ Study Guide], Chapter 10: Working with Git, Section: Ignoring Files with .gitignore
? [How to Use .gitignore File]

NEW QUESTION 80
A Linux administrator cloned an existing Linux server and built a new server from that clone. The administrator encountered the following error after booting the cloned server:

```
Device mismatch detected
```

The administrator performed the commands listed below to further troubleshoot and mount the missing filesystem:

```
#ls -al /dev/disk/by-uuid/
total 0
drwxr-xr-x 2 root 220 Jul 08:59 .
drwxr-xr-x 2 root 160 Jul 08:59 ..
lrwxrwxrwx 1 root 26 Jul 11:10 2251a54-6c14-9187-df8629373 -> ../../sdb
lrwxrwxrwx 1 root 26 Jul 11:10 4211c54-2a13-7291-bd8629373 -> ../../sdc
lrwxrwxrwx 1 root 26 Jul 11:10 3451b54-6d10-3561-ad8629373 -> ../../sdd
```

Which of the following should administrator use to resolve the device mismatch issue and mount the disk?

A. mount disk by device-id
B. fsck -A
C. mount disk by-label
D. mount disk by-blkid

Answer: A

Explanation:
The administrator should use the command mount disk by device-id to resolve the device mismatch issue and mount the disk. The issue is caused by the cloned server having a different device name for the disk than the original server. The output of blkid shows that the disk has the device name /dev/sdb1 on the cloned server, but the output of cat /etc/fstab shows that the disk is expected to have the device name /dev/sda1. The command mount disk by device-id will mount the disk by using its unique identifier (UUID) instead of its device name. The UUID can be obtained from the output of blkid or lsblk -f. The command will mount the disk to the specified mount point (/data) and resolve the issue. The other options are incorrect because they either do not mount the disk (fsck -A), do not use the correct identifier (mount disk by-label or mount disk by-blkid), or do not exist (mount disk by-blkid). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 318-319.

NEW QUESTION 83
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 85
A systems administrator needs to check if the service systemd-resolved.service is running without any errors. Which of the following commands will show this information?

A. systemct1 status systemd-resolved.service
B. systemct1 enable systemd-resolved.service
C. systemct1 mask systemd-resolved.service
D. systemct1 show systemd-resolved.service

Answer: A

Explanation:
The command systemct1 status systemd-resolved.service will show the information about the service systemd-resolved.service. The systemct1 command is a tool for managing system services and units. The status option displays the current status of a unit, such as active, inactive, or failed. The output also shows the unit description, loaded configuration, process ID, memory usage, and recent log messages. This command will show if the service systemd-resolved.service is running without any errors. This is the correct command to use to accomplish the task. The other options are incorrect because they either perform different actions (enable, mask, or show) or do not show the status of the service (systemct1 show systemd-resolved.service only shows the properties of the service, not the status). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, page 427.

NEW QUESTION 89
A systems administrator is tasked with installing GRUB on the legacy MBR of the SATA hard drive. Which of the following commands will help the administrator accomplish this task?

A. grub-install /dev/hda
B. grub-install /dev/sda
C. grub-install /dev/sr0
D. grub-install /dev/hd0,0

Answer: B

Explanation:
The command that will help the administrator install GRUB on the legacy MBR of the SATA hard drive is grub-install /dev/sda. This command will install GRUB on the master boot record (MBR) of the first SATA disk (/dev/sda). The MBR is the first sector of a disk that contains boot code and a partition table. GRUB will overwrite the boot code and place its own code that can load GRUB modules and configuration files from a specific partition.
The other options are not correct commands for installing GRUB on the legacy MBR of the SATA hard drive. The grub-install /dev/hda command will try to install GRUB on the first IDE disk (/dev/hda), which may not exist or may not be bootable. The grub-install /dev/sr0 command will try to install GRUB on the first SCSI CD-ROM device (/dev/sr0), which is not a hard drive and may not be bootable. The grub-install /dev/hd0,0 command is invalid because grub-install does not accept partition names as arguments, only disk names. References: Installing GRUB using grub-install; GRUB Manual

NEW QUESTION 93
A junior systems administrator recently installed an HBA card in one of the servers that is deployed for a production environment. Which of the following commands can the administrator use to confirm on which server the card was installed?

A. lspci | egrep 'hba| fibr'
B. lspci | zgrep 'hba | fibr'
C. lspci | pgrep 'hba| fibr'
D. lspci | 'hba | fibr'

Answer: A

Explanation:
The best command to use to confirm on which server the HBA card was installed is A. lspci | egrep ‘hba| fibr’. This command will list all the PCI devices on the server and filter the output for those that match the pattern ‘hba’ or ‘fibr’, which are likely to be related to the HBA card. The egrep command is a variant of grep that supports extended regular expressions, which allow the use of the ‘|’ operator for alternation. The other commands are either invalid or will not produce the desired output. For example:
? B. lspci | zgrep ‘hba | fibr’ will try to use zgrep, which is a command for searching compressed files, not standard output.
? C. lspci | pgrep ‘hba| fibr’ will try to use pgrep, which is a command for finding processes by name or other attributes, not text patterns.
? D. lspci | ‘hba | fibr’ will try to use ‘hba | fibr’ as a command, which is not valid and will cause an error.

NEW QUESTION 96
A Linux administrator is adding a new configuration file to a Git repository. Which of the following describes the correct order of Git commands to accomplish the task successfully?

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

Answer: B

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

NEW QUESTION 99
A systems administrator is adding a Linux-based server and removing a Windows-based server from a cloud-based environment. The changes need to be validated before they are applied to the cloud-based environment. Which of the following tools should be used to meet this requirement?

A. Ansible
B. git clone
C. git pull
D. terraform plan

Answer: D

Explanation:
Terraform is a tool for building, changing, and managing infrastructure as code in a cloud- based environment. Terraform uses configuration files to describe the desired state of the infrastructure and applies changes accordingly. Terraform supports various cloud providers, such as AWS, Azure, Google Cloud Platform, and more.
To validate changes before they are applied to the cloud-based environment, the administrator can use the terraform plan command. This command will compare the current state of the infrastructure with the desired state defined in the configuration files and show what actions will be performed to achieve the desired state. This command will not make any changes to the infrastructure but only show a plan of changes. The statement D is correct.
The statements A, B, and C are incorrect because they do not validate changes before they are applied to the cloud-based environment. Ansible is another tool for automating infrastructure management, but it does not have a plan command. Git clone and git pull are commands for working with git repositories, which are used for version control of code. References: [How to Use Terraform to Manage Cloud Infrastructure]

NEW QUESTION 102
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 104
Which of the following will prevent non-root SSH access to a Linux server?

A. Creating the /etc/nologin file
B. Creating the /etc/nologin.allow file containing only a single line root
C. Creating the /etc/nologin/login.deny file containing a single line +all
D. Ensuring that /etc/pam.d/sshd includes account sufficient pam_nologin.so

Answer: A

Explanation:
This file prevents any non-root user from logging in to the system, regardless of the authentication method. The contents of the file are displayed to the user before the login is terminated. This can be useful for system maintenance or security reasons12.
References: 1: Creating the /etc/nologin File - Oracle 2: How to Restrict Log In Capabilities of Users on Ubuntu

NEW QUESTION 109
A Linux engineer has been notified about the possible deletion of logs from the file
/opt/app/logs. The engineer needs to ensure the log file can only be written into without removing previous entries.

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

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

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

Answer: A

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

NEW QUESTION 112
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 113
Users are unable to create new files on the company's FTP server, and an administrator is troubleshooting the issue. The administrator runs the following commands:

```
# df -h /ftpusers/
Filesystem    Size     Used     Avail    Use%    Mounted on
/dev/sda4     150G     40G      109G     26%     /ftpusers

# df -i /ftpusers/

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

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

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

Answer: C

Explanation:
The cause of the issue based on the output above is C. The inodes is at full capacity and would affect file creation for users.
An inode is a data structure that stores information about a file or directory, such as its name, size, permissions, owner, timestamps, and location on the disk. Each file or directory has a unique inode number that identifies it. The number of inodes on a filesystem is fixed when the filesystem is created, and it determines how many files and directories can be created on that filesystem. If the inodes are exhausted, no new files or directories can be created, even if there is enough disk space available.
The output for the second command shows that the /ftpusers/ filesystem has 0% of inodes available, which means that all the inodes have been used up. This

would prevent users from creating new files on the FTP server. The administrator should either delete some unused files or directories to free up some inodes, or resize the filesystem to increase the number of inodes.
The other options are incorrect because:
* A. The users do not have the correct permissions to create files on the FTP server.
This is not true, because the output for the first command shows that the /ftpusers/ filesystem has 26% of disk space available, which means that there is enough space for users to create files. The permissions of the files and directories are not shown in the output, but they are not relevant to the issue of inode exhaustion.
* B. The ftpusers filesystem does not have enough space.
This is not true, because the output for the first command shows that the /ftpusers/ filesystem has 26% of disk space available, which means that there is enough space for users to create files. The issue is not related to disk space, but to inode capacity.
* D. ftpusers is mounted as read only.
This is not true, because the output for the first command does not show any indication that the /ftpusers/ filesystem is mounted as read only. If it was, it would have an (ro) flag next to the mounted on column. A read only filesystem would prevent users from creating or modifying files on the FTP server, but it would not affect the inode usage.

NEW QUESTION 117
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 it12. References: 1: Git - git-fetch Documentation 2: Git Fetch | Atlassian Git Tutorial

NEW QUESTION 122
A systems administrator was tasked with assigning the temporary IP address/netmask 192.168.168.1/255.255.255.255 to the interface eth0 of a Linux server. When adding the address, the following error appears:
ip address add 192.168.168.1/33 dev eth0
Error: any valid prefix is expected rather than "192.168.168.1/33".
Based on the command and its output above, which of the following is the cause of the issue?

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

Answer: A

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

NEW QUESTION 127
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 128
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 131
A Linux administrator is troubleshooting an issue in which an application service failed to start on a Linux server. The administrator runs a few commands and gets the following outputs:

Output 1:

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

Output 2:

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

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

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

Answer: B

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

NEW QUESTION 133
A Linux administrator found many containers in an exited state. Which of the following commands will allow the administrator to clean up the containers in an exited state?

A. docker rm --all
B. docker rm $(docker ps -aq)
C. docker images prune *
D. docker rm --state exited

Answer: B

Explanation:
The command docker rm $(docker ps -aq) will allow the administrator to clean up the containers in an exited state. The docker command is a tool for managing Docker containers on Linux systems. Docker containers are isolated and lightweight environments that can run applications and services without affecting the host system. Docker uses images to create containers, which are files that contain the code, libraries, dependencies, and configuration of the applications and services. The rm option removes one or more containers. The $(docker ps -aq) is a command substitution that executes the command inside the parentheses and replaces it with the output. The docker ps - aq command lists all the containers, including the ones in an exited state, and shows only their IDs. The docker rm $(docker ps -aq) command will remove all the containers, including the ones in an exited state, by passing their IDs to the rm option. This will allow the administrator to clean up the containers in an exited state. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not exist (docker rm --all or docker rm --state exited) or do not remove the containers (docker images prune *). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 571.

NEW QUESTION 136
A Linux engineer needs to download a ZIP file and wants to set the nice of value to -10 for this new process. Which of the following commands will help to accomplish the task?

A. $ nice -v -10 wget https://foo.com/installation.zip
B. $ renice -v -10 wget https://foo.com/installation.2ip
C. $ renice -10 wget https://foo.com/installation.zip
D. $ nice -10 wget https://foo.com/installation.zip

Answer: D

Explanation:
The nice -10 wget https://foo.com/installation.zip command will help to accomplish the task of downloading a ZIP file and setting the nice value to -10 for this new process. The nice command can be used to run a program with a modified scheduling priority, which affects how much CPU time the process receives. The nice value ranges from -20 (highest priority) to 19 (lowest priority), and the default value is 0. The -10 option specifies the nice value to be used for the wget command, which will download the ZIP file from the given URL. The nice -v -10 wget https://foo.com/installation.zip command is incorrect, as -v is not a valid option for nice. The renice -v -10 wget https://foo.com/installation.zip command is incorrect, as renice is used to change the priority of an existing process, not a new one. The renice -10 wget https://foo.com/installation.zip command is incorrect for the same reason as above. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Memory and Process Execution, page 469.

NEW QUESTION 139
A Linux system is having issues. Given the following outputs:
dig @192.168.2.2 mycomptiahost
; << >> 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.References1: How To Troubleshoot DNS Client Issues in Linux - RootUsers2: 6 Best Tools to Troubleshoot DNS
Issues in Linux - Tecmint3: How To Troubleshoot DNS in Linux - OrcaCore4: Fixing DNS Issues in Ubuntu 20.04 | DeviceTests

NEW QUESTION 140
A systems administrator needs to remove a disk from a Linux server. The disk size is 500G, and it is the only one that size on that machine. Which of the following commands can the administrator use to find the corresponding device name?

A. fdisk -V
B. partprobe -a
C. lsusb -t
D. lsscsi -s

Answer: D

Explanation:
The lsscsi command can list the SCSI devices on the system, along with their size and device name. The -s option shows the size of each device. The administrator can look for the device that has a size of 500G and note its device name. See lsscsi(8) - Linux man page and How to check Disk Interface Types in Linux.References1: https://linux.die.net/man/8/lsscsi2: https://www.golinuxcloud.com/check-disk-type-linux/

NEW QUESTION 141
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 144
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 146
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 150
A Linux administrator has defined a systemd script docker-repository.mount to mount a volume for use by the Docker service. The administrator wants to ensure that Docker service does not start until the volume is mounted. Which of the following configurations needs to be added to the Docker service definition to best accomplish this task?

A. After=docker-respository.mount
B. ExecStart=/usr/bin/mount -a
C. Requires=docker-repository.mount
D. RequiresMountsFor=docker-repository.mount

Answer: C

Explanation:
This option declares an explicit dependency between the Docker service and the docker- repository.mount unit. It means that the Docker service will not start unless the docker- repository.mount unit is successfully activated. This ensures that the volume is mounted before the Docker service tries to use it12.
References: 1: systemd.unit - systemd unit configuration 2: How to mount host volumes into docker containers in Dockerfile during build

NEW QUESTION 153
Which of the following specifications is used to perform disk encryption in a Linux system?

A. LUKS
B. TLS
C. SSL
D. NFS

Answer: A

Explanation:
LUKS stands for Linux Unified Key Setup, which is a specification for disk encryption on Linux systems. LUKS allows users to encrypt partitions or entire disks using a passphrase or a key file. LUKS also supports multiple keys and key slots, which can be used to unlock the encrypted data. LUKS is compatible with various tools and utilities, such as cryptsetup, dm-crypt, and LVM. References: [How to Encrypt Partitions with LUKS on Linux]

NEW QUESTION 154
A Linux administrator is trying to start the database service on a Linux server but is not able to run it. The administrator executes a few commands and receives the following output:

```
#systemctl status mariadb
mariadb.servcice
   Loaded: masked (Reason: Unit mariadb.service is masked)
   Active: inactive (dead)

#systemctl enable mariadb
Failed to enable unit: ...

#systemctl start mariadb
Failed to start mariadb.service ...
```

Which of the following should the administrator run to resolve this issue? (Select two).

A. systemctl unmask mariadb
B. journalctl —g mariadb
C. dnf reinstall mariadb
D. systemctl start mariadb
E. chkconfig mariadb on
F. service mariadb reload

Answer: AD

Explanation:
These commands will unmask the mariadb service, which is currently prevented from starting, and then start it normally. The other commands are either not relevant, not valid, or not sufficient for this task. For more information on how to manage masked services with systemctl, you can refer to the web search result 1.

NEW QUESTION 157
An administrator created an initial Git repository and uploaded the first files. The administrator sees the following when listing the repository:

```
__init__.py      Initial Commit    Just now
main.py          Initial Commit    Just now
.DS_STORE        Initial Commit    Just now
setup.sh         Initial Commit    Just now
README.md        Initial Commit    Just now
```

The administrator notices the file . DS STORE should not be included and deletes it from the online repository. Which of the following should the administrator run from the root of the local repository before the next commit to ensure the file is not uploaded again in future commits?

A. rm -f .DS STORE && git push
B. git fetch && git checkout .DS STORE
C. rm -f .DS STORE && git rebase origin main
D. echo .DS STORE >> .gitignore

Answer: D

Explanation:
The correct answer is D. The administrator should run “echo .DS STORE >> .gitignore” from the root of the local repository before the next commit to ensure the file is not uploaded again in future commits.
This command will append the file name .DS STORE to the end of the .gitignore file, which is a special file that tells Git to ignore certain files or directories that should not be tracked or uploaded to the repository. By adding .DS STORE to the .gitignore file, the administrator will prevent Git from staging, committing, or pushing this file in the future.
The other options are incorrect because:
* A. rm -f .DS STORE && git push
This command will delete the file .DS STORE from the local repository and then push the changes to the remote repository. However, this does not prevent the file from being uploaded again in future commits, if it is recreated or copied to the local repository.
* B. git fetch && git checkout .DS STORE
This command will fetch the latest changes from the remote repository and then restore the file .DS STORE from the remote repository to the local repository. This is not what the administrator wants to do, as this will undo the deletion of the file from the online repository.
* C. rm -f .DS STORE && git rebase origin main
This command will delete the file .DS STORE from the local repository and then rebase the local branch onto the main branch of the remote repository. This will rewrite the commit history of the local branch and may cause conflicts or errors. This is not what the administrator wants to do, as this is a risky and unnecessary operation.

NEW QUESTION 162
A Linux system is failing to boot. The following error is displayed in the serial console: [[1;33mDEPEND[Om] Dependency failed for /data.
[[1;33mDEPEND[Om] Dependency failed for Local File Systems
...
Welcome to emergency mode! After logging in, type "journalctl -xb" to viewsystem logs,
"systemct1 reboot" to reboot, "systemct1 default" to try again to boot into default mode.
Give root password for maintenance (or type Control-D to continue}
Which of the following files will need to be modified for this server to be able to boot again?

A. /etc/mtab
B. /dev/sda

C. /etc/fstab
D. /ete/grub.conf

Answer: C

Explanation:
The file that will need to be modified for the server to be able to boot again is /etc/fstab. The /etc/fstab file is a file that contains the information about the file systems that are mounted at boot time on Linux systems. The file specifies the device name, mount point, file system type, mount options, dump frequency, and pass number for each file system. The error message indicates that the dependency failed for /data, which is a mount point for a file system. This means that the system could not mount the /data file system at boot time, which caused the system to enter the emergency mode. The emergency mode is a mode that allows the administrator to log in as the root user and perform basic tasks such as repairing the system. The administrator should modify the /etc/fstab file and check the entry for the /data file system. The administrator should look for any errors or inconsistencies in the device name, file system type, or mount options, and correct them. The administrator should also verify that the device and the file system are intact and functional by using commands such as blkid, fdisk, fsck, or mount. The administrator should then reboot the system and see if the issue is resolved. The file that will need to be modified for the server to be able to boot again is /etc/fstab. This is the correct answer to the question. The other options are incorrect because they are not related to the file systems that are mounted at boot time (/etc/mtab, /dev/sda,
or /etc/grub.conf). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 321.

NEW QUESTION 163
A Linux administrator is configuring a two-node cluster and needs to be able to connect the nodes to each other using SSH keys from the root account. Which of the following commands will accomplish this task?

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

Answer: C

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

NEW QUESTION 165
Which of the following would significantly help to reduce data loss if more than one drive fails at the same time?

A. Server clustering
B. Load balancing
C. RAID
D. VDI

Answer: C

Explanation:
RAID stands for Redundant Array of Independent Disks, which is a technology that combines multiple physical disks into a logical unit that provides improved performance, reliability, or both. RAID can significantly help to reduce data loss if more than one drive fails at the same time, depending on the RAID level used. For example, RAID 1 (mirroring) duplicates the data on two or more disks, so that if one disk fails, the data can be recovered from another disk. RAID 5 (striping with parity) distributes the data and parity information across three or more disks, so that if one disk fails, the data can be reconstructed from the remaining disks. RAID 6 (striping with double parity) extends RAID 5 by adding another parity block, so that if two disks fail, the data can still be recovered from the remaining disks.
References: [What is RAID?]

NEW QUESTION 170
Which of the following technologies can be used as a central repository of Linux users and groups?

A. LDAP
B. MFA
C. SSO
D. PAM

Answer: A

Explanation:
LDAP stands for Lightweight Directory Access Protocol, which is a protocol for accessing and managing a central directory of users and groups. LDAP can be used as a central repository of Linux users and groups, allowing for centralized authentication and authorization across multiple Linux systems. MFA, SSO, and PAM are not technologies that can be used as a central repository of Linux users and groups. MFA stands for Multi- Factor Authentication, which is a method of verifying a user's identity using more than one factor, such as a password, a token, or a biometric. SSO stands for Single Sign-On, which is a feature that allows a user to log in once and access multiple applications or systems without having to re-enter credentials. PAM stands for Pluggable Authentication Modules, which is a framework that allows Linux to use different authentication methods, such as passwords, tokens, or biometrics. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing Users and Groups

NEW QUESTION 174
A systems administrator needs to verify whether the built container has the app.go file in its root directory. Which of the following can the administrator use to verify the root directory has this file?

A. docker image inspect
B. docker container inspect
C. docker exec <container_name> ls
D. docker ps <container_name>

Answer: C

Explanation:
The docker exec <container_name> ls command can be used to verify whether the built container has the app.go file in its root directory. This command will run the ls command inside the specified container and list the files and directories in its root directory. If the app.go file is present, it will be displayed in the output. The docker image inspect command will display information about an image, not a container, and it will not list the files inside the image. The docker container inspect command will display information about a container, not its files. The docker ps <container_name> command is invalid, as ps does not accept a container name as an argument. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Virtualization and Cloud Technologies, page 499.

NEW QUESTION 175
A systems administrator has been tasked with disabling the nginx service from the environment to prevent it from being automatically and manually started. Which of the following commands will accomplish this task?

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

Answer: C

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

NEW QUESTION 180
A Linux engineer receives reports that files created within a certain group are being modified by users who are not group members. The engineer wants to reconfigure the server so that only file owners and group members can modify new files by default. Which of the following commands would accomplish this task?

A. chmod 775
B. umas
C. 002
D. chactr -Rv
E. chown -cf

Answer: B

Explanation:
The command umask 002 will accomplish the task of reconfiguring the server so that only file owners and group members can modify new files by default. The umask command is a tool for setting the default permissions for new files and directories on Linux systems. The umask value is a four-digit octal number that represents the permissions that are subtracted from the default permissions. The default permissions for files are 666, which means read and write for owner, group, and others. The default permissions for directories are 777, which means read, write, and execute for owner, group, and others. The umask value consists of four digits: the first digit is for special permissions, such as setuid, setgid, and sticky bit; the second digit is for the owner permissions; the third digit is for the group permissions; and the fourth digit is for the others permissions. The umask value can be calculated by subtracting the desired permissions from the default permissions. For example, if the desired permissions for files are 664, which means read and write for owner and group, and read for others, then the umask value is 002, which is 666 - 664. The command umask 002 will set the umask value to 002, which will ensure that only file owners and group members can modify new files by default. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not set the default permissions for new files (chmod 775 or chown - cf) or do not exist (chattr -Rv). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing File Permissions and Ownership, page 349.

NEW QUESTION 185
Several users reported that they were unable to write data to the /oracle1 directory. The following output has been provided:

Filesystem	Size	Used	Available	Use%	Mounted on
/dev/sdb1	100G	50G	50G	50%	/oracle1

Which of the following commands should the administrator use to diagnose the issue?

A. df -i /oracle1
B. fdisk -1 /dev/sdb1
C. lsblk /dev/sdb1
D. du -sh /oracle1

Answer: A

Explanation:
The administrator should use the command df -i /oracle1 to diagnose the issue of users being unable to write data to the /oracle1 directory. This command will show the inode usage of the /oracle1 filesystem, which indicates how many files and directories can be created on it. If the inode usage is 100%, it means that no more files or directories can be added, even if there is still free space on the disk. The administrator can then delete some unnecessary files or directories, or increase the inode limit of the filesystem, to resolve the issue.
The other options are not correct commands for diagnosing this issue. The fdisk -l /dev/sdb1 command will show the partition table of /dev/sdb1, which is not relevant to the inode usage. The lsblk /dev/sdb1 command will show information about /dev/sdb1 as a block device, such as its size, mount point, and type, but not

its inode usage. The du -sh /oracle1 command will show the disk usage of /oracle1 in human-readable format, but not its inode usage. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 7: Managing Disk Storage; How to Check Inode Usage in Linux - Fedingo

NEW QUESTION 186
A cloud engineer needs to remove all dangling images and delete all the images that do not have an associated container. Which of the following commands will help to accomplish this task?

A. docker images prune -a
B. docker push images -a
C. docker rmi -a images
D. docker images rmi --all

Answer: A

Explanation:
The command docker images prune -a will help to remove all dangling images and delete all the images that do not have an associated container.
The docker command is a tool for managing Docker containers and images.
The images subcommand operates on images. The prune option removes unused images.
The -a option removes all images, not just dangling ones. A dangling image is an image that is not tagged and is not referenced by any container. This command will accomplish the task of cleaning up the unused images. The other options are incorrect because they either do not exist (docker push images -a or docker images rmi --all) or do not remove images (docker rmi -a images only removes images that match the name or ID of “images”). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 567.

NEW QUESTION 189
A new disk was presented to a server as /dev/ sdd. The systems administrator needs to check if a partition table is on that disk. Which of the following commands can show this information?

A. lsscsi
B. fdisk
C. blkid
D. partprobe

Answer: B

Explanation:
The command that can be used to check if a partition table is on a disk is fdisk. The fdisk command can display, create, delete, and modify partitions on a disk. To show the partition table of a disk, the administrator can use fdisk -l /dev/sdd (B). References:
? [CompTIA Linux+ Study Guide], Chapter 5: Managing Filesystems and Logical Volumes, Section: Partitioning Disks
? [How to Use Fdisk Command in Linux]

NEW QUESTION 190
A systems administrator is troubleshooting connectivity issues and trying to find out why a Linux server is not able to reach other servers on the same subnet it is connected to. When listing link parameters, the following is presented:

```
# ip link list dev eth0
2: etho: <NO-CARRIER, BROADCAST, MULTICAST, UP> mtu 1500, qdisc
fq_codel state DOWN mode DEFAULT group default qlen 1000
link/ether ac:00:11:22:33:cd brd ff:ff:ff:ff:ff:ff
```

Based on the output above, which of following is the MOST probable cause of the issue?

A. The address ac:00:11:22:33:cd is not a valid Ethernet address.
B. The Ethernet broadcast address should be ac:00:11:22:33:ff instead.
C. The network interface eth0 is using an old kernel module.
D. The network interface cable is not connected to a switch.

Answer: D

Explanation:
The most probable cause of the connectivity issue is that the network interface cable is not connected to a switch. This can be inferred from the output of the ip link list dev eth0 command, which shows that the network interface eth0 has the NO- CARRIER flag set. This flag indicates that there is no physical link detected on the interface, meaning that the cable is either unplugged or faulty. The other options are not valid causes of the issue. The address ac:00:11:22:33:cd is a valid Ethernet address, as it follows the format of six hexadecimal octets separated by colons. The Ethernet broadcast address should be ff:ff:ff:ff:ff:ff, which is the default value for all interfaces. The network interface eth0 is not using an old kernel module, as it shows the UP flag, which indicates that the interface is enabled and ready to transmit data. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Networking

NEW QUESTION 193
Which of the following will prevent non-root SSH access to a Linux server?

A. Creating the /etc/nologin file
B. Creating the /etc/nologin.allow file containing only a single line root
C. Creating the /etc/nologin/login.deny file containing a single line +all
D. Ensuring that /etc/pam.d/sshd includes account sufficient pam_nologin.so

Answer: A

Explanation:

This file prevents any non-root user from logging in to the system, regardless of the authentication method. The contents of the file are displayed to the user before the login is terminated. This can be useful for system maintenance or security reasons12.
References: 1: Creating the /etc/nologin File - Oracle 2: How to Restrict Log In Capabilities of Users on Ubuntu

NEW QUESTION 195
Which of the following commands is used to configure the default permissions for new files?

A. setenforce
B. sudo
C. umask
D. chmod

Answer: C

Explanation:
The command that is used to configure the default permissions for new files is umask. The umask command is a tool for setting the default permissions for new files and directories on Linux systems. The umask value is a four-digit octal number that represents the permissions that are subtracted from the default permissions. The default permissions for files are 666, which means read and write for owner, group, and others.
The default permissions for directories are 777, which means read, write, and execute for owner, group, and others. The umask value consists of four digits: the first digit is for special permissions, such as setuid, setgid, and sticky bit; the second digit is for the owner permissions; the third digit is for the group permissions; and the fourth digit is for the others permissions. The umask value can be calculated by subtracting the desired permissions from the default permissions. For example, if the desired permissions for files are 664, which means read and write for owner and group, and read for others, then the umask value is 002, which is 666 - 664. The command umask 002 will set the umask value to 002, which will ensure that only file owners and group members can modify new files by default. The command that is used to configure the default permissions for new files is umask. This is the correct answer to the question. The other options are incorrect because they either do not set the default permissions for new files (setenforce, sudo, or chmod) or do not exist (kill -HUP or kill -TERM).
References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing File Permissions and Ownership, page 349.

NEW QUESTION 200
A systems administrator wants to test the route between IP address 10.0.2.15 and IP address 192.168.1.40. Which of the following commands will accomplish this task?

A. route -e get to 192.168.1.40 from 10.0.2.15
B. ip route get 192.163.1.40 from 10.0.2.15
C. ip route 192.169.1.40 to 10.0.2.15
D. route -n 192.168.1.40 from 10.0.2.15

Answer: B

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

NEW QUESTION 202
A Linux engineer finds multiple failed login entries in the security log file for application users. The Linux engineer performs a security audit and discovers a security issue. Given the following:
grep -iE '*www*|db' /etc/passwd
www-data:x:502:502:www-data:/var/www:/bin/bash db:x: 505:505:db: /opt/db:/bin/bash
Which of the following commands would resolve the security issue?

A. usermod -d /srv/www-data www-data && usermod -d /var/lib/db db
B. passwd -u www-data && passwd -u db
C. renice -n 1002 -u 502 && renice -n 1005 -u 505
D. chsh -s /bin/false www-data && chsh -s /bin/false db

Answer: D

Explanation:
This command will use the chsh tool to change the login shell of the users www-data and db to /bin/false, which means they will not be able to log in to the system1. This will prevent unauthorized access attempts and improve security.
References: 1: Replacing /bin/bash with /bin/false in /etc/passwd file

NEW QUESTION 206
A systems administrator is troubleshooting a connectivity issue pertaining to access to a system named db.example.com. The system IP address should be 192.168.20.88. The administrator issues the dig command and receives the following output:

```
;; ANSWER SECTION:
db.example.com.     15 IN A 192.168.20.89
```

The administrator runs grep db.example.com /etc/hosts and receives the following output:

```
192.168.20.89 db.example.com
```

Given this scenario, which of the following should the administrator do to address this issue?

A. Modify the /etc/hosts file and change the db.example.com entry to 192.168.20.89.

B. Modify the /etc/network file and change the db.example.com entry to 192.168.20.88.
C. Modify the /etc/network file and change the db.example.com entry to 192.168.20.89.
D. Modify the /etc/hosts file and change the db.example.com entry to 192.168.20.88.

Answer: D

Explanation:
The administrator should modify the /etc/hosts file and change the db.example.com entry to 192.168.20.88 to address the issue. The /etc/hosts file is a file that maps hostnames to IP addresses on Linux systems. The file can be used to override the DNS resolution and provide a local lookup for hostnames. The dig output shows that the DNS returns the IP address 192.168.20.88 for the hostname db.example.com, which is the correct IP address of the system. The grep output shows that the /etc/hosts file contains an entry for db.example.com with the IP address 192.168.20.89, which is the wrong IP address of the system. This can cause a conflict and prevent the system from being accessed by the hostname. The administrator should modify the /etc/hosts file and change the db.example.com entry to 192.168.20.88, which is the correct IP address of the system. This will align the /etc/hosts file with the DNS and allow the system to be accessed by the hostname. The administrator should modify the /etc/hosts file and change the db.example.com entry to 192.168.20.88 to address the issue. This is the correct answer to the question. The other options are incorrect because they either do not modify the /etc/hosts file (modify the /etc/network file and change the db.example.com entry to 192.168.20.88 or modify the /etc/network file and change the db.example.com entry to 192.168.20.89) or do not change the IP address to the correct one (modify the /etc/hosts file and change the db.example.com entry to 192.168.20.89). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 378.

NEW QUESTION 207
A Linux administrator needs to analyze a failing application that is running inside a container. Which of the following commands allows the Linux administrator to enter the running container and analyze the logs that are stored inside?

A. docker run -ti app /bin/sh
B. podman exec -ti app /bin/sh
C. podman run -d app /bin/bash
D. docker exec -d app /bin/bash

Answer: B

Explanation:
Podman exec -ti app /bin/sh allows the Linux administrator to enter the running container and analyze the logs that are stored inside. This command uses the podman tool, which is a daemonless container engine that can run and manage containers on Linux systems. The exec option executes a command inside an existing container, in this case app, which is the name of the container that runs the failing application. The -ti option allocates a pseudo-TTY and keeps STDIN open, allowing for interactive shell access to the container. The /bin/sh argument specifies the shell command to run inside the container, which can be used to view and manipulate the log files.
The other options are not correct commands for entering a running container and analyzing the logs. Docker run -ti app /bin/sh creates a new container from the app image and runs the /bin/sh command inside it, but does not enter the existing container that runs the failing application. Podman run -d app /bin/bash also creates a new container from the app image and runs the /bin/bash command inside it, but does so in detached mode, meaning that it runs in the background without interactive shell access. Docker exec -d app /bin/bash executes the /bin/bash command inside the existing app container, but also does so in detached mode, without interactive shell access.
References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks; View container logs | Docker Docs; How to see the logs of a docker container - Stack Overflow

NEW QUESTION 211
A developer is trying to install an application remotely that requires a graphical interface for installation. The developer requested assistance to set up the necessary environment variables along with X11 forwarding in SSH. Which of the following environment variables must be set in remote shell in order to launch the graphical interface?

A. $RHOST
B. SETENV
C. $SHELL
D. $DISPLAY

Answer: D

Explanation:
The environment variable that must be set in remote shell in order to launch the graphical interface is $DISPLAY. This variable tells X11 applications where to display their windows on screen. It usually has the form hostname:displaynumber.screennumber, where hostname is the name of the computer running the X server, displaynumber is a unique identifier for an X display on that computer, and screennumber is an optional identifier for a screen within an X display. For example, localhost:0.0 means display number 0 on the local host. If the hostname is omitted, it defaults to the local host.
The other options are not correct environment variables for launching the graphical interface. $RHOST is a variable that stores the name of the remote host, but it is not used by X11 applications. SETENV is a command that sets environment variables in some shells, but it is not an environment variable itself. $SHELL is a variable that stores the name of the current shell, but it is not related to X11 forwarding. References: How to enable or disable X11 forwarding in an SSH server; How to Configure X11 Forwarding Using SSH In Linux

NEW QUESTION 213
A Linux administrator needs to create a new user named user02. However, user02 must be in a different home directory, which is under /comptia/projects. Which of the following commands will accomplish this task?

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

Answer: A

Explanation:
The command useradd -d /comptia/projects user02 will accomplish the task of creating a new user named user02 with a different home directory.
The useradd command is a tool for creating new user accounts on Linux systems. The - d option specifies the home directory for the new user, which is the

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

NEW QUESTION 216
Ann, a security administrator, is performing home directory audits on a Linux server. Ann issues the su Joe command and then issues the Is command. The output displays files that reside in Ann's home directory instead of Joe's. Which of the following represents the command Ann should have issued in order to list Joe's files?

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

Answer: A

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

NEW QUESTION 221
A user is attempting to log in to a Linux server that has Kerberos SSO ena-bled. Which of the following commands should the user run to authenticate and then show the ticket grants? (Select TWO).

A. kinit
B. klist
C. kexec
D. kioad
E. pkexec
F. realm

Answer: AB

Explanation:
The following commands can help the user to authenticate and show the ticket grants using Kerberos SSO on a Linux server:
? kinit: This command obtains and caches an initial ticket-granting ticket (TGT) for
the user from the Kerberos key distribution center (KDC). The user needs to enter their password or use a keytab file to authenticate1.
? klist: This command lists the cached tickets, including the TGT and any service tickets, for the user. It also shows the expiration time and flags for each ticket2.
For example, the user can run the following commands to log in and view their tickets:
$ kinit username@REALM Password for username@REALM:
$ klist
Ticket cache: FILE:/tmp/krb5cc_1000 Default principal: username@REALM
Valid starting Expires Service principal
04/06/2023 16:06:59 04/07/2023 02:06:59 krbtgt/REALM@REALM
renew until 04/13/2023 16:06:59 References:
? kinit(1) - Linux man page, section "Description".
? klist(1) - Linux man page, section "Description".

NEW QUESTION 226
Which of the following tools is BEST suited to orchestrate a large number of containers across many different servers?

A. Kubernetes
B. Ansible
C. Podman
D. Terraform

Answer: A

Explanation:
The tool that is best suited to orchestrate a large number of containers across many different servers is Kubernetes. Kubernetes is an open-source platform for managing containerized applications and services. Kubernetes allows the administrator to deploy, scale, and update containers across a cluster of servers, as well as to automate the configuration and coordination of the containers. Kubernetes also provides features such as service discovery, load balancing, storage management, security, monitoring, and logging. Kubernetes can handle complex and dynamic workloads and ensure high availability and performance of the containers. Kubernetes is the tool that is best suited to orchestrate a large number of containers across many different servers. This is the correct answer to the question. The other options are incorrect because they either do not orchestrate containers (Ansible or Terraform) or do not operate across many different servers (Podman). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 573.

NEW QUESTION 227
A systems engineer is adding a new 1GB XFS filesystem that should be temporarily mounted under /ops/app. Which of the following is the correct list of commands to achieve this goal?

A.

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

B.

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

C.

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

D.

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

Answer: D

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

NEW QUESTION 232
A Linux administrator was tasked with deleting all files and directories with names that are contained in the sobelete.txt file. Which of the following commands will accomplish this task?

A. xargs -f cat toDelete.txt -rm
B. rm -d -r -f toDelete.txt
C. cat toDelete.txt | rm -frd
D. cat toDelete.txt | xargs rm -rf

Answer: D

Explanation:
The command cat toDelete.txt | xargs rm -rf will delete all files and directories with names that are contained in the toDelete.txt file. The cat command reads the file and outputs its contents to the standard output. The | operator pipes the output to the next command. The xargs command converts the output into arguments for the next command. The rm -rf command removes the files and directories recursively and forcefully. This is the correct way to accomplish the task. The other options are incorrect because they either use the wrong options (-f instead of -a for xargs), the wrong arguments (toDelete.txt instead of toDelete.txt filename for rm), or the wrong commands (rm instead of xargs). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, pages 349-350.

NEW QUESTION 236
SIMULATION
Junior system administrator had trouble installing and running an Apache web server on a Linux server. You have been tasked with installing the Apache web server on the Linux server and resolving the issue that prevented the junior administrator from running Apache.
INSTRUCTIONS
Install Apache and start the service. Verify that the Apache service is running with the defaults.
Typing “help” in the terminal will show a list of relevant event commands.
If at any time you would like to bring back the initial state of the simulation, please click the Reset All button.

CentOS Command Prompt

A. Mastered
B. Not Mastered

Answer: A

Explanation:
yum install httpd
systemct1 --now enable httpd systemct1 status httpd netstat -tunlp | grep 80
pkill <processname> systemct1 restart httpd systemct1 status httpd

NEW QUESTION 240
A Linux administrator would like to use systemd to schedule a job to run every two hours. The administrator creates timer and service definitions and restarts the server to load these new configurations. After the restart, the administrator checks the log file and notices that the job is only running daily. Which of the following is MOST likely causing the issue?

A. The checkdiskspace.service is not running.
B. The checkdiskspace.service needs to be enabled.
C. The OnCalendar schedule is incorrect in the timer definition.
D. The system-daemon services need to be reloaded.

Answer: C

Explanation:
The OnCalendar schedule is incorrect in the timer definition, which is causing the issue. The OnCalendar schedule defines when the timer should trigger the service. The format of the schedule is OnCalendar=<year>-<month>-<day> <hour>:<minute>:<second>. If any of the fields are omitted, they are assumed to be *, which means any value. Therefore, the schedule OnCalendar=*-*-* 00:00:00 means every day at midnight, which is why the job is running daily. To make the job run every two hours, the schedule should be OnCalendar=*-*-* *:00:00/2, which means every hour divisible by 2 at the start of the minute. The other options are incorrect because they are not related to the schedule. The checkdiskspace.service is running, as shown by the output of systemct1 status checkdiskspace.service. The checkdiskspace.service is enabled, as shown by the output of systemct1 is-enabled checkdiskspace.service. The system-daemon services do not need to be reloaded, as the timer and service definitions are already loaded by the
restart. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Processes and Scheduling Tasks, page 437.

NEW QUESTION 244
A systems administrator created a new Docker image called test. After building the image, the administrator forgot to version the release. Which of the following will allow the administrator to assign the v1 version to the image?

A. docker image save test test:v1
B. docker image build test:vl
C. docker image tag test test:vl
D. docker image version test:v1

Answer: C

Explanation:
The docker image tag test test:v1 command can be used to assign the v1 version to the image called test. This command creates a new tag for the existing image, without changing the original image. The docker image save test test:v1 command would save the image to a file, not assign a version. The docker image build test:vl command is invalid, as vl is not a valid version number. The docker image version test:v1 command does not exist. References: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 16: Virtualization and Cloud Technologies, page 500.

NEW QUESTION 245
A server is experiencing intermittent connection issues. Some connections to the Internet work as intended, but some fail as if there is no connectivity. The systems administrator inspects the server configuration:

Routing table:

```
default via 89.107.157.129 dev ens3 proto static metric 100
default via 10.0.5.1 dev ens11 proto dhcp metric 101
10.0.0.0/16 dev sn11 proto kernel scope link src 10.0.6.225 metric 101
89.107.157.128/26 via 89.107.157.129 dev ens3 proto static metric 100
89.107.157.129 dev ens3 proto static scope link metric 100
89.107.157.160/29 dev ens3 proto kernel scope link src 89.107.157.161 metric 100
```

IP configuration:

```
ens3:
  inet 89.107.157.161/29 brd 89.107.157.167 scope global neprefixroute ens3
ens11:
  inet 10.0.6.225/16 brd 10.0.255.255 scope global noprefixroute dynamic ens11
```

ARP table:

Address	Hwtype	Hwaddress	Flags	Mask	Iface
10.0.5.1	ether	64:d1:54:c4:75:cb	C		ens11
89.107.157.129	ether	5c:5e:ab:01:85:cf	C		ens3
89.107.157.162	ether	52:54:00:e1:44:0a	C		ens3
10.0.255.1	ether	00:50:7f:e3:aa:1c	C		ens11

/etc/resolv.conf:

```
Generated by NetworkHanager
search company.com
nameserver 10.0.5.1
```

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

A. An internal-only DNS server is configured.
B. The IP netmask is wrong for ens3.
C. Two default routes are configured.
D. The ARP table contains incorrect entries.

Answer: C

Explanation:
The most likely cause of the issue is that two default routes are configured on the server. The default route is the route that is used when no other route matches the destination of a packet. The default route is usually the gateway that connects the local network to the Internet. The server configuration shows that there are two default routes in the routing table, one with the gateway 192.168.1.1 and the other with the gateway 10.0.0.1. This can cause a conflict and confusion for the server when deciding which gateway to use for the outgoing packets. Some packets may be sent to the wrong gateway and fail to reach the Internet, while some packets may be sent to the correct gateway and work as intended. This can result in intermittent connection issues and inconsistent behavior. The administrator should remove one of the default routes and keep only the correct one for the network. This can be done by using the ip route del command or by editing the network configuration files. This will resolve the issue and restore the connectivity. The other options are incorrect because they are not supported by the outputs. The DNS server, the IP netmask, and the ARP table are not the causes of the issue. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, pages 381-382.

NEW QUESTION 250
A systems administrator created a web server for the company and is required to add a tag for the API so end users can connect. Which of the following would the administrator do to complete this requirement?

A. hostnamectl status --no-ask-password
B. hostnamectl set-hostname "$(perl -le "print" "A" x 86)"
C. hostnamectl set-hostname Comptia-WebNode -H root@192.168.2.14
D. hostnamectl set-hostname Comptia-WebNode --transient

Answer: C

Explanation:
The command hostnamectl set-hostname Comptia-WebNode -H root@192.168.2.14 sets the hostname of the web server to Comptia-WebNode and connects to the server using the SSH protocol and the root user. This is the correct way to complete the requirement. The other options are incorrect because they either display the current hostname status (hostnamectl status), set an invalid hostname (hostnamectl set- hostname "$(perl -le "print" "A" x 86)"), or set a transient hostname that is not persistent (hostnamectl set-hostname Comptia-WebNode --transient). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 9: Managing System Components, page 291.

NEW QUESTION 252
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

XK0-005 Practice Exam Features:

* XK0-005 Questions and Answers Updated Frequently

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

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

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