CompTIA

Exam Questions XK0-005

CompTIA Linux+ Certification Exam

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1
A Linux administrator was notified that a virtual server has an I/O bottleneck. The Linux administrator analyzes the following output:

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

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

root@linux:~# free -m

              total  used    free  shared buff/cache  available
Mem:           3933  1454     110      33        2368       2202
Swap:          1497     5    1491
```

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

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

Answer: B

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

NEW QUESTION 2
A Linux administrator intends to start using KVM on a Linux server. Which of the following commands will allow the administrator to load the KVM module as well as any related dependencies?

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

Answer: A

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

NEW QUESTION 3
A Linux administrator needs to remove software from the server. Which of the following RPM options should be used?

A. rpm -s
B. rm -d
C. rpm -q
D. rpm -e

Answer: D

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

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

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

Answer: A

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

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

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

Answer: C

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

NEW QUESTION 6
A Linux administrator has physically added a new RAID adapter to a system. Which of the following commands should the Linux administrator run to confirm that the device has been recognized? (Select TWO).

A. rmmod
B. Is -11 /etc
C. Ishw —class disk
D. pvdisplay
E. rmdir /dev
F. dmesg

Answer: CF

Explanation:
The following commands can help you confirm that the new RAID adapter has been recognized by the Linux system:
? dmesg: This command displays the kernel messages, which can show the information about the newly detected hardware device. You can use dmesg | grep -i raid to filter the output for RAID-related messages.
? lshw -class disk: This command lists the disk devices on the system, including the RAID controller and its model name. You can use lshw -class disk | grep -i raid to filter the output for RAID-related information1.
The other commands are not relevant for this purpose. For example:
? rmmod: This command removes a module from the Linux kernel, which is not useful for detecting a new device.
? ls -l /etc: This command lists the files and directories in the /etc directory, which is not related to hardware devices.
? pvdisplay: This command displays the attributes of physical volumes, which are part of the logical volume management (LVM) system, not the RAID system.
? rmdir /dev: This command removes an empty directory, which is not helpful for detecting a new device. Moreover, /dev is a special directory that contains device files, and should not be removed.

NEW QUESTION 7
A Linux administrator is configuring a new internal web server fleet. The web servers are up and running but can only be reached by users directly via IP address. The administrator is attempting to fix this inconvenience by requesting appropriate records from the DNS team. The details are:
Hostname: devel.comptia.org
IP address: 5.5.5.1, 5.5.5.2, 5.5.5.3, 5.5.5.4
Name server: 5.5.5.254
Additional names: dev.comptia.org, development.comptia.org
Which of the following types of DNS records should the Linux administrator request from the DNS team? (Select three).

A. MX
B. NS
C. PTR

D. A
E. CNAME
F. RRSIG
G. SOA
H. TXT
I. SRV

Answer: BDE

Explanation:
The Linux administrator should request the following types of DNS records from the DNS team:
? A: This record type is used to map a hostname to an IPv4 address. The administrator needs four A records for devel.comptia.org, one for each IP address (5.5.5.1, 5.5.5.2, 5.5.5.3, 5.5.5.4). This will allow users to access the web servers by using the hostname devel.comptia.org instead of the IP addresses1.
? CNAME: This record type is used to create an alias for another hostname. The administrator needs two CNAME records, one for dev.comptia.org and one for development.comptia.org, both pointing to devel.comptia.org. This will allow users to access the web servers by using any of these three hostnames interchangeably1.
? NS: This record type is used to delegate a domain or a subdomain to another name server. The administrator needs one NS record for comptia.org, pointing to 5.5.5.254, which is the name server that hosts the records for the subdomain devel.comptia.org2. This will allow users to resolve the hostnames under comptia.org by querying the name server 5.5.5.2542.
The other record types are not relevant for the administrator's task:
? MX: This record type is used to specify the mail exchange server for a domain or a subdomain1. The administrator does not need this record type because the web servers are not intended to handle email traffic.
? PTR: This record type is used to map an IP address to a hostname, which is the reverse of an A record1. The administrator does not need this record type because the web servers are not expected to be accessed by their IP addresses.
? RRSIG: This record type is used to provide digital signatures for DNSSEC, which is a security extension for DNS that verifies the authenticity and integrity of DNS responses3. The administrator does not need this record type because it is not mentioned in the task requirements.
? SOA: This record type is used to provide information about the authoritative name server and other parameters for a domain or a subdomain1. The administrator does not need this record type because it is usually created automatically by the name server software when a new zone file is created4.
? TXT: This record type is used to store arbitrary text data that can be used for various purposes, such as SPF, DKIM, DMARC, etc1. The administrator does not need this record type because it is not related to the web server functionality.
? SRV: This record type is used to specify the location and port number of a service that runs on a domain or a subdomain1. The administrator does not need this record type because the web servers use the standard HTTP port 80, which does not require an SRV record.
References: 1: DNS Record Types – CompTIA Network+ N10-007 – 1.8 2: NS Record - DNSimple Help 3: DNSSEC - Wikipedia 4: SOA Record - DNSimple Help

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 Linux administrator is troubleshooting a systemd mount unit file that is not working correctly. The file contains:
[root@system] # cat mydocs.mount [Unit]
Description=Mount point for My Documents drive [Mount]
What=/dev/drv/disk/by-uuid/94afc9b2-ac34-ccff-88ae-297ab3c7ff34 Where=/home/user1/My Documents
Options=defaults Type=xfs
[Install]
WantedBy=multi-user.target
The administrator verifies the drive UUID correct, and user1 confirms the drive should be mounted as My Documents in the home directory. Which of the following can the administrator do to fix the issues with mounting the drive? (Select two).

A. Rename the mount file to home-user1-My\x20Documents.mount.

B. Rename the mount file to home-user1-my-documents.mount.
C. Change the What entry to /dev/drv/disk/by-uuid/94afc9b2\-ac34\-ccff\-88ae\- 297ab3c7ff34.
D. Change the Where entry to Where=/home/user1/my\ documents.
E. Change the Where entry to Where=/home/user1/My\x20Documents.
F. Add quotes to the What and Where entries, such as What="/dev/drv/disk/by- uuid/94afc9b2-ac34-ccff-88ae-297ab3c7ff34" and Where="/home/user1/My Documents".

Answer: AE

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

NEW QUESTION 10
A cloud engineer needs to block the IP address 192.168.10.50 from accessing a Linux server. Which of the following commands will achieve this goal?

A. iptables -F INPUT -j 192.168.10.50 -m DROP
B. iptables -A INPUT -s 192.168.10.30 -j DROP
C. iptables -i INPUT --ipv4 192.168.10.50 -z DROP
D. iptables -j INPUT 192.168.10.50 -p DROP

Answer: B

Explanation:
The correct command to block the IP address 192.168.10.50 from accessing a Linux server is iptables -A INPUT -s 192.168.10.50 -j DROP. This command appends a rule to the INPUT chain that matches the source address 192.168.10.50 and jumps to the DROP target, which discards the packet. The other commands are incorrect because they either have invalid syntax, wrong parameters, or wrong order of arguments. References:
CompTIA Linux+ Study Guide, Fourth Edition, page 457-458.

NEW QUESTION 10
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 12
The application team has reported latency issues that are causing the application to crash on the Linux server. The Linux administrator starts troubleshooting and receives the following output:

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

# ethtool -S eth0
rx_fw_discards: 4487
```

Which of the following commands will improve the latency issue?

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

Answer: D

Explanation:
The best command to use to improve the latency issue is D. # echo ‘net.core.rmem max = 12500000’ >> /etc/sysctl.conf # echo ‘net.core.wmem_max = 12500000’ >> /etc/sysctl.conf # sysctl -p. This command will increase the size of the receive and send buffers for the network interface, which can improve the network performance and reduce packet loss. The sysctl command will apply the changes to the kernel parameters without rebooting the system.
The other commands are either incorrect or not suitable for this task. For example:

? A. # echo 'net.core.net_backlog = 5000000' >> /etc/sysctl.conf # sysctl -p # systemctl daemon-reload will try to increase the backlog queue for incoming connections, but this is not relevant for the latency issue. The systemctl daemon- reload command is also unnecessary, as it only reloads the systemd configuration files, not the kernel parameters.
? B. # ifdown eth0 # ip link set dev eth0 mtu 800 # ifup eth0 will try to change the maximum transmission unit (MTU) of the network interface to 800 bytes, but this is too low and may cause fragmentation and performance degradation. The default MTU for Ethernet is 1500 bytes, and it should not be changed unless there is a specific reason.
? C. # systemctl stop network # ethtool -g eth0 512 # systemctl start network will try to change the ring buffer size of the network interface to 512, but this is too small and may cause packet drops and latency spikes. The default ring buffer size for Ethernet is usually 4096 or higher, and it should be increased if there is a high network traffic.

NEW QUESTION 14
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:xx SRC=x.x.x.x DST=x.x.x.x LEN=1059 TOS=0x00
PREC=0x00 TTL=115 ID=31368 DF PROTO=TCP
SPT=17992 DPT=80 WINDOW=16477 RES=0x00 ACK PSH URGP=0
Oct 20 03:46:02 hostname kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx SRC=x.x.x.x DST=x.x.x.x LEN=52 TOS=0x00
PREC=0x00 TTL=52 ID=763 DF PROTO=TCP SPT=20229 DPT=22 WINDOW=15598 RES=0x00 ACK URGP=0
Oct 20 03:46:14 hostname kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx SRC=x.x.x.x DST=x.x.x.x LEN=324 TOS=0x00
PREC=0x00 TTL=49 ID=64245 PROTO=TCP SPT=47237 DPT=80 WINDOW=470 RES=0x00 ACK PSH URGP=0
Oct 20 03:46:26 hostname kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx 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 15
To harden one of the servers, an administrator needs to remove the possibility of remote administrative login via the SSH service. Which of the following should the administrator do?

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

Answer: B

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

NEW QUESTION 17
An administrator needs to make an application change via a script that must be run only in console mode. Which of the following best represents the sequence the administrator should execute to accomplish this task?

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

Answer: A

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

NEW QUESTION 18
User1 is a member of the accounting group. Members of this group need to be able to execute but not make changes to a script maintained by User2. The script should not be accessible to other users or groups. Which of the following will give proper access to the script?

A. chown user2:accounting script.sh chmod 750 script.sh
B. chown user1:accounting script.shchmod 777 script.sh
C. chown accounting:user1 script.sh chmod 057 script.sh
D. chown user2:accounting script.sh chmod u+x script.sh

Answer: A

Explanation:
The commands that will give proper access to the script are:
? chown user2:accounting script.sh: This command will change the ownership of the script to user2 as the owner and accounting as the group. The chown command is a tool for changing the owner and group of files and directories on Linux systems. The user2:accounting is the user and group name that the command should assign to the script. The script.sh is the name of the script that the command should modify. The command chown user2:accounting script.sh will ensure that user2 is the owner of the script and accounting is the group of the script, which will allow user2 to maintain the script and the accounting group to access the script.
? chmod 750 script.sh: This command will change the permissions of the script to 750, which means read, write, and execute for the owner; read and execute for the group; and no access for others. The chmod command is a tool for changing the permissions of files and directories on Linux systems. The permissions are represented by three digits in octal notation, where each digit corresponds to the owner, group, and others. Each digit can have a value from 0 to 7, where each value represents a combination of read, write, and execute permissions. The 750 is the permission value that the command should assign to the script.
The script.sh is the name of the script that the command should modify. The command chmod 750 script.sh will ensure that only the owner and the group can execute the script, but not make changes to it, and that the script is not accessible to other users or groups.
The commands that will give proper access to the script are chown user2:accounting script.sh and chmod 750 script.sh. This is the correct answer to the question.
The other options are incorrect because they either do not give proper access to the script (chown user1:accounting script.sh or chown accounting:user1 script.sh) or do not change the permissions of the script (chmod 777 script.sh or chmod u+x
script.sh). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing File Permissions and Ownership, pages 346-348.

NEW QUESTION 21
Rugged appliances are small appliances with ruggedized hardware and like Quantum Spark appliance they use which operating system?

A. Centos Linux
B. Gaia embedded
C. Gaia
D. Red Hat Enterprise Linux version 5

Answer: B

Explanation:
Rugged appliances are small appliances with ruggedized hardware that use Gaia embedded as their operating system. Gaia embedded is a version of Gaia that is optimized for embedded devices such as Rugged appliances and Quantum Spark appliances. Gaia embedded supports features such as VPN, firewall, identity awareness, application control, URL filtering, and anti-bot. Gaia embedded does not use Centos Linux, Gaia, or Red Hat Enterprise Linux version 5 as their

operating system. References: Check Point Rugged Appliance Datasheet, page 1.

NEW QUESTION 26
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 29
A user generated a pair of private-public keys on a workstation. Which of the following commands will allow the user to upload the public key to a remote server and enable passwordless login?

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

Answer: D

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

NEW QUESTION 30
A systems administrator requires that all files that are created by the user named web have read-only permissions by the owner. Which of the following commands will satisfy this requirement?

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

Answer: C

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

NEW QUESTION 35
Users are reporting that writes on a system configured with SSD drives have been taking longer than expected, but reads do not seem to be affected. A Linux systems administrator is investigating this issue and working on a solution. Which of the following should the administrator do to help solve the issue?

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

Answer: A

Explanation:
TRIM is a feature that allows the operating system to inform the SSD which blocks of data are no longer in use and can be wiped internally. This helps to maintain the SSD’s performance and endurance by preventing unnecessary write operations and reducing write amplification12. Running the corresponding command to trim the SSD drives, such as fstrim or blkdiscard on Linux, can help to solve the issue of slow writes by freeing up space and optimizing the SSD’s internal garbage collection34.
References: 1: What is SSD TRIM, why is it useful, and how to check whether it is turned on 2: How to Trim SSD in Windows 10 3: How to run fsck on an external

drive with OS X? 4: How to Use the fsck Command on Linux

NEW QUESTION 36
A Linux administrator is installing a web server and needs to check whether web traffic has already been allowed through the firewall. Which of the following commands should the administrator use to accomplish this task?

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

Answer: C

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

NEW QUESTION 39
A systems administrator wants to be sure the sudo rules just added to /etc/sudoers are valid. Which of the following commands can be used for this task?

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

Answer: A

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

NEW QUESTION 44
The development team wants to prevent a file from being modified by all users in a Linux system, including the root account. Which of the following commands can be used to accomplish this objective?

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

Answer: C

Explanation:
The chattr command is used to change file attributes on Linux systems that support extended attributes, such as ext2, ext3, ext4, btrfs, xfs, and others. File attributes are flags that modify the behavior of files and directories.
To prevent a file from being modified by all users in a Linux system, including the root account, the development team can use the chattr +i /app/conf/file command. This command will set the immutable attribute (+i) on the file /app/conf/file, which means that the file cannot be deleted, renamed, linked, appended, or written to by any user or process. To remove the immutable attribute, the development team can use the chattr -i /app/conf/file command. The statement C is correct.
The statements A, B, and D are incorrect because they do not prevent the file from being modified by all users. The chmod /app/conf/file command does not work because it requires an argument to specify the permissions to change. The setenforce /app/conf/file command does not work because it is used to change the SELinux mode, not file attributes. The chmod 0000 /app/conf/file command will remove all permissions from the file, but it can still be modified by the root account.
References: [How to Use chattr Command in Linux]

NEW QUESTION 47
A systems administrator is gathering information about a file type and the contents of a file. Which of the following commands should the administrator use to accomplish this task?

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

Answer: A

Explanation:

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

NEW QUESTION 49
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 54
An engineer needs to insert a character at the end of the current line in the vi text editor. Which of the following will allow the engineer to complete this task?

A. p
B. r
C. bb
D. A
E. i

Answer: D

Explanation:
The vi text editor is a popular and powerful tool for editing text files on Linux systems. The vi editor has two modes: command mode and insert mode. In command mode, the user can issue commands to manipulate the text, such as moving the cursor, deleting, copying, pasting, searching, replacing, and saving. In insert mode, the user can type text into the file. To switch from command mode to insert mode, the user can press various keys, such as i, a, o, I, A, or O. To switch from insert mode to command mode, the user can press the Esc key.
To insert a character at the end of the current line in the vi editor, the user can press the A key in command mode. This will move the cursor to the end of the line and switch to insert mode. Then, the user can type the desired character and press Esc to return to command mode. The statement D is correct.
The statements A, B, C, and E are incorrect because they do not perform the desired task. The p key in command mode will paste the previously copied or deleted text after the cursor. The r key in command mode will replace the character under the cursor with another character. The bb key in command mode will move the cursor back two words. The i key in command mode will switch to insert mode before the cursor. References: [How to Use vi Text Editor in Linux]

NEW QUESTION 57
A DevOps engineer wants to allow the same Kubernetes container configurations to be deployed in development, testing, and production environments. A key requirement is that the containers should be configured so that developers do not have to statically configure custom, environment-specific locations. Which of the following should the engineer use to meet this requirement?

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

Answer: D

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

NEW QUESTION 61
A DevOps engineer needs to download a Git repository from https://git.company.com/admin/project.git. Which of the following commands will achieve this goal?

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

Answer: A

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

Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 571.

NEW QUESTION 62
While inspecting a recently compromised Linux system, the administrator identified a number of processes that should not have been running:

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

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

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

Answer: A

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

NEW QUESTION 65
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 70
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 72
A Linux systems administrator receives reports from various users that an application hosted on a server has stopped responding at similar times for several days in a row. The administrator logs in to the system and obtains the following output:
Output 1:

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

Output 2:

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

Output 3:

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

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

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

Answer: D

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

NEW QUESTION 74
A Linux system is getting an error indicating the root filesystem is full. Which of the following commands should be used by the systems administrator to resolve this issue? (Choose three.)

A. df -h /
B. fdisk -1 /dev/sdb
C. growpart /dev/mapper/rootvg-rootlv
D. pvcreate /dev/sdb
E. lvresize –L +10G -r /dev/mapper/rootvg-rootlv
F. lsblk /dev/sda
G. parted -l /dev/mapper/rootvg-rootlv
H. vgextend /dev/rootvg /dev/sdb

Answer: ACE

Explanation:
The administrator should use the following three commands to resolve the issue of the root filesystem being full:
? df -h /. This command will show the disk usage of the root filesystem in a human- readable format. The df command is a tool for reporting file system disk space usage. The -h option displays the sizes in powers of 1024 (e.g., 1K, 234M, 2G). The / specifies the root filesystem. The command df -h / will show the total size, used space, available space, and percentage of the root filesystem. This command will help the administrator identify the problem and plan the solution.
? growpart /dev/mapper/rootvg-rootlv. This command will grow the partition that contains the root filesystem to the maximum size available.
The growpart command is a tool for resizing partitions on Linux systems. The /dev/mapper/rootvg-rootlv is the device name of the partition, which is a logical volume managed by the Logical Volume Manager (LVM). The command growpart /dev/mapper/rootvg-rootlv will extend the partition to fill the disk space and increase the size of the root filesystem. This command will help the administrator solve the problem and free up space.
? lvresize –L +10G -r /dev/mapper/rootvg-rootlv. This command will resize the logical volume that contains the root filesystem and add 10 GB of space.
The lvresize command is a tool for resizing logical volumes on Linux systems. The -L option specifies the new size of the logical volume, in this case +10G, which means 10 GB more than the current size. The -r option resizes the underlying file system as well. The /dev/mapper/rootvg-rootlv is the device name of the logical volume, which is the same as the partition name. The command lvresize –L +10G -r /dev/mapper/rootvg-rootlv will increase the size of the logical volume and the root filesystem by 10 GB and free up space. This command will help the administrator solve the problem and free up space.
The other options are incorrect because they either do not affect the root filesystem (fdisk -1 /dev/sdb, pvcreate /dev/sdb, lsblk /dev/sda, or vgextend /dev/rootvg /dev/sdb) or do not use the correct syntax (fdisk -1 /dev/sdb instead of fdisk -l /dev/sdb or parted -l /dev/mapper/rootvg-rootlv instead of parted /dev/mapper/rootvg-rootlv print). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 318-319, 331-332.

NEW QUESTION 78
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 82
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 85
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 86
Which of the following is the best tool for dynamic tuning of kernel parameters?

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

Answer: A

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

NEW QUESTION 91
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 95
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 98
A Linux administrator wants to set the SUID of a file named dev_team.text with 744 access rights. Which of the following commands will achieve this goal?

A. chmod 4744 dev_team.txt
B. chmod 744 --setuid dev_team.txt
C. chmod -c 744 dev_team.txt
D. chmod -v 4744 --suid dev_team.txt

Answer: A

Explanation:
The command that will set the SUID of a file named dev_team.txt with 744 access rights is chmod 4744 dev_team.txt. This command will use the chmod utility to change the file mode bits of dev_team.txt. The first digit (4) represents the SUID bit, which means that when someone executes dev_team.txt, it will run with the permissions of the file owner. The next three digits (744) represent the read, write, and execute permissions for the owner (7), group (4), and others (4). This means that the owner can read, write, and execute dev_team.txt, while the group and others can only read it.
The other options are not correct commands for setting the SUID of a file with 744 access rights. The chmod 744 --setuid dev_team.txt command is invalid because there is no -- setuid option in chmod. The chmod -c 744 dev_team.txt command will change the file mode bits to 744, but it will not set the SUID bit. The -c option only means that chmod will report when a change is made. The chmod -v 4744 --suid dev_team.txt command is also invalid because there is no --suid option in chmod. The -v option only means that chmod will output a diagnostic for every file processed. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing Users and Groups; chmod(1) - Linux manual page

NEW QUESTION 99
A junior administrator is setting up a new Linux server that is intended to be used as a router at a remote site. Which of the following parameters will accomplish this goal?
A.

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

A.

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

B.

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

C.

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

Answer: C

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

NEW QUESTION 100
Which of the following actions are considered good security practices when hardening a Linux server? (Select two).

A. Renaming the root account to something else
B. Removing unnecessary packages
C. Changing the default shell to /bin/csh
D. Disabling public key authentication
E. Disabling the SSH root login possibility
F. Changing the permissions on the root filesystem to 600

Answer: BE

Explanation:
Some good security practices when hardening a Linux server are:
? Removing unnecessary packages (B) to reduce the attack surface and eliminate potential vulnerabilities
? Disabling the SSH root login possibility (E) to prevent unauthorized access and brute-force attacks on the root account References:
? [CompTIA Linux+ Study Guide], Chapter 9: Securing Linux, Section: Hardening Linux
? [How to Harden Your Linux Server]

NEW QUESTION 103
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 105
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 108
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 113
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 117
A systems administrator is receiving tickets from users who cannot reach the application app that should be listening on port 9443/tcp on a Linux server.
To troubleshoot the issue, the systems administrator runs netstat and receives the following output:

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

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

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

Answer: B

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

NEW QUESTION 122
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 123
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 128
An administrator would like to securely connect to a server and forward port 8080 on a local machine to port 80 on the server. Which of the following commands should the administrator use to satisfy both requirements?

A. ssh —L 8080: localhost:80 admin@server
B. ssh —R 8080: localhost:80 admin@server
C. ssh —L 80 : localhost:8080 admin@server
D. ssh —R 80 : localhost:8080 admin@server

Answer: A

Explanation:
This command will create a local port forwarding, which means that connections from the SSH client are forwarded via the SSH server, then to a destination server. In this case, the destination server is the same as the SSH server (localhost), and the destination port is 80. The SSH client will listen on port 8080 on the local machine, and any connection to that port will be forwarded to port 80 on the server. This way, the administrator can securely access the web service running on port 80 on the server by using http://localhost:8080 on the local machine.
The other options are incorrect because:
* B. ssh -R 8080:localhost:80 admin@server
This command will create a remote port forwarding, which means that connections from the SSH server are forwarded via the SSH client, then to a destination server. In this case, the destination server is the same as the SSH client (localhost), and the destination port is 80. The SSH server will listen on port 8080 on the remote machine, and any connection to that port will be forwarded to port 80 on the client. This is not what the administrator wants to do.
* C. ssh -L 80:localhost:8080 admin@server
This command will also create a local port forwarding, but it will use port 80 on the local machine and port 8080 on the server. This is not what the administrator wants to do, and it may also fail if port 80 is already in use by another service on the local machine.
* D. ssh -R admin@server
This command is incomplete and invalid. It does not specify any port numbers or destination addresses for the remote port forwarding. It will also fail if the SSH server does not allow remote port forwarding.
References:
? CompTIA Linux+ Certification Exam Objectives
? How to Set up SSH Tunneling (Port Forwarding)

NEW QUESTION 133
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 134
A systems administrator is compiling a report containing information about processes that are listening on the network ports of a Linux server. Which of the following commands will allow the administrator to obtain the needed information?

A. ss -pint
B. tcpdump -nL
C. netstat -pn
D. lsof -It

Answer: A

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

NEW QUESTION 137
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 140
A Linux systems administrator is setting up a new web server and getting 404 - NOT FOUND errors while trying to access the web server pages from the browser. While working on the diagnosis of this issue, the Linux systems administrator executes the following commands:

```
# getenforce
Enforcing

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

Which of the following commands will BEST resolve this issue?

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

Answer: B

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

NEW QUESTION 142
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 146
DRAG DROP
A new drive was recently added to a Linux system. Using the environment and tokens provided, complete the following tasks:
• Create an appropriate device label.
• Format and create an ext4 file system on the new partition. The current working directory is /.

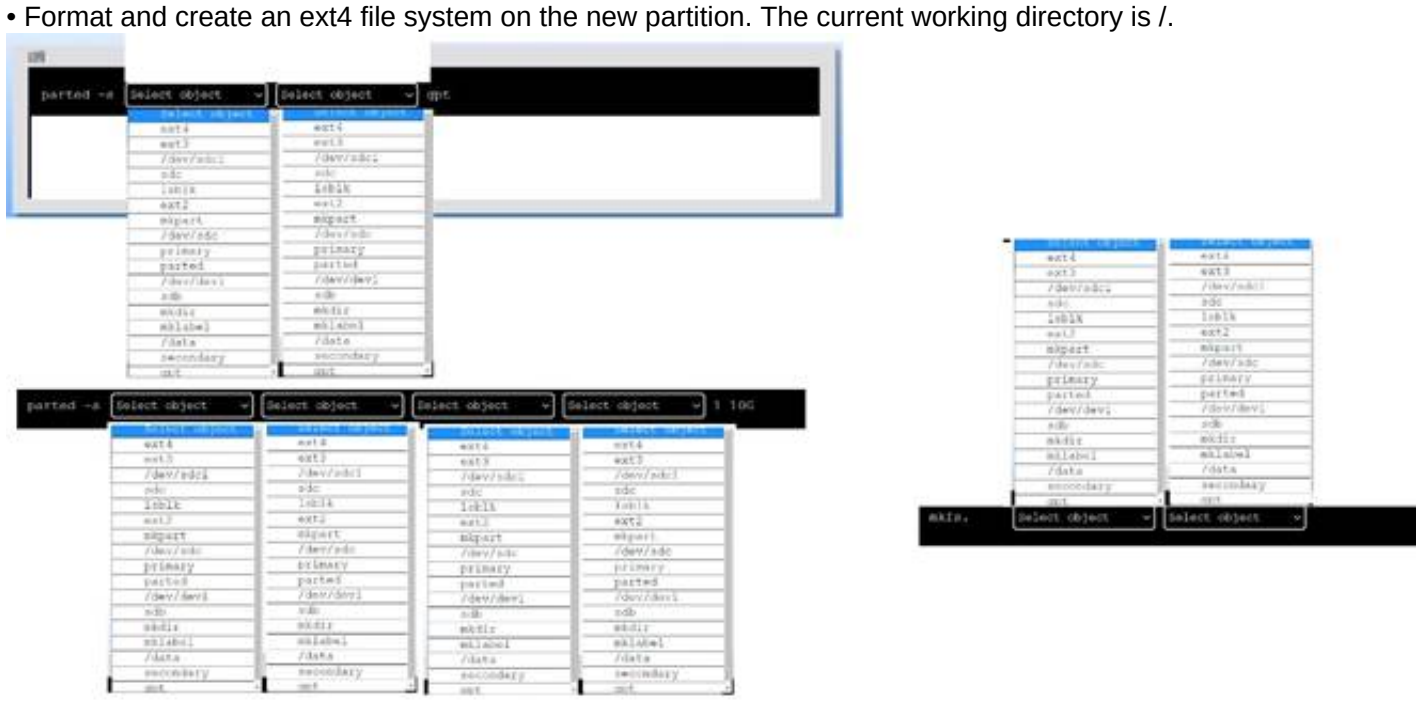

A. Mastered
B. Not Mastered

Answer: A

Explanation:
To create an appropriate device label, format and create an ext4 file system on the new partition, you can use the following commands:
? To create a GPT (GUID Partition Table) label on the new drive /dev/sdc, you can use the parted command with the -s option (for script mode), the device name (/dev/sdc), the mklabel command, and the label type (gpt). The command is:
parted -s /dev/sdc mklabel gpt
? To create a primary partition of 10 GB on the new drive /dev/sdc, you can use the parted command with the -s option, the device name (/dev/sdc), the mkpart command, the partition type (primary), the file system type (ext4), and the start and end points of the partition (1 and 10G). The command is:
parted -s /dev/sdc mkpart primary ext4 1 10G
? To format and create an ext4 file system on the new partition /dev/sdc1, you can use the mkfs command with the file system type (ext4) and the device name (/dev/sdc1). The command is:
mkfs.ext4 /dev/sdc1
You can verify that the new partition and file system have been created by using the lsblk command, which will list all block devices and their properties.

NEW QUESTION 148
A user created the following script file:
! /bin/bash
FILENAME: /home/user/ script . sh echo "hello world"
exit 1
However, when the user tried to run the script file using the command "script . sh, an error returned indicating permission was denied. Which of the follow-ing should the user execute in order for the script to run properly?

A. chmod u+x /home/user/script . sh
B. chmod 600 /home/user/script . sh
C. chmod /home/user/script . sh
D. chmod 0+r /horne/user/scrip

E. sh

Answer: A

Explanation:
To run a script file, the user needs to have execute permission on the file. The command chmod u+x /home/user/script.sh (A) will grant execute permission to the owner of the file, which is the user who created it. The other commands will not give execute permission to the user, and therefore will not allow the script to run properly. References:
? [CompTIA Linux+ Study Guide], Chapter 3: Working with Files, Section: Changing File Permissions
? [How to Make a Bash Script Executable]

NEW QUESTION 150
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 155
An administrator attempts to connect to a remote server by running the following command:
$ nmap 192.168.10.36
Starting Nmap 7.60 (https://nmap.org) at 2022-03-29 20:20 UTC Nmap scan report for www1 (192.168.10.36)
Host is up (0.000091s latency). Not shown: 979 closed ports PORT STATE SERVICE 21/tcp open ftp 22/tcp filtered ssh 631/tcp open ipp
Nmap done: 1 IP address (1 host up) scanned in 0.06 seconds
Which of the following can be said about the remote server?

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

Answer: A

Explanation:
This is because the port 22/tcp is shown as filtered by nmap, which means that nmap cannot determine whether the port is open or closed because a firewall or other device is blocking its probes. If the SSH server was not running on the remote server, the port would be shown as closed, which means that nmap received a TCP RST packet in response to its probe. If the remote SSH server was using SSH protocol version 1, the port would be shown as open, which means that nmap received a TCP SYN/ACK packet in response to its probe. If the SSH host key on the remote server had expired, the port would also be shown as open, but the SSH client would display a warning message about the host key verification failure. Therefore, the best explanation for the filtered state of the port 22/tcp is that a firewall is preventing nmap from reaching the SSH server.
You can find more information about nmap port states and how to interpret them in the following web search results:
? Nmap scan what does STATE=filtered mean?
? How to find ports marked as filtered by nmap
? Technical Tip: NMAP scan shows ports as filtered

NEW QUESTION 158
An administrator deployed a Linux server that is running a web application on port 6379/tcp.
SELinux is in enforcing mode based on organization policies. The port is open on the firewall.
Users who are trying to connect to a local instance of the web application receive Error 13, Permission denied.
The administrator ran some commands that resulted in the following output:

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

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

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

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

Answer: B

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

NEW QUESTION 159
A systems administrator is encountering performance issues. The administrator runs 3 commands with the following output

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

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

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

Answer: A

Explanation:
Based on the output of the image sent by the user, the system is under CPU pressure and will require additional vCPUs. The output shows that there are four processes running upload.sh scripts that are consuming a high percentage of CPU time (99.7%, 99.6%, 99.5%, and 99.4%). The output also shows that the system has only 4 vCPUs, which means that each process is using almost one entire vCPU. This indicates that the system is struggling to handle the CPU load and may experience performance issues or slowdowns. Adding more vCPUs to the system would help to alleviate the CPU pressure and improve the system performance. The system has not been running for over a year, as the uptime command shows that it has been up for only 1 day, 2 hours, and 13 minutes. The number of users logged in to the system is not relevant to the performance issue, as they are not consuming significant CPU resources. The system does not require more memory, as the free command shows that it has plenty of available memory (49 GB total, 48 GB free). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Memory and Process Execution, pages 468-469.

NEW QUESTION 160
A Linux administrator needs to redirect all HTTP traffic temporarily to the new proxy server 192.0.2.25 on port 3128. Which of the following commands will accomplish this task?

A. iptables -t nat -D PREROUTING -p tcp --sport 80 -j DNAT - -to-destination 192.0.2.25:3128
B. iptables -t nat -A PREROUTING -p top --dport 81 -j DNAT –-to-destination 192.0.2.25:3129
C. iptables -t nat -I PREROUTING -p top --sport 80 -j DNAT –-to-destination 192.0.2.25:3129
D. iptables -t nat -A PREROUTING -p tcp --dport 80 -j DNAT –-to-destination 192.0.2.25:3128

Answer: D

Explanation:
The command iptables -t nat -A PREROUTING -p tcp --dport 80 -j DNAT -- to-destination 192.0.2.25:3128 adds a rule to the nat table that redirects all incoming TCP packets with destination port 80 (HTTP) to the proxy server 192.0.2.25 on port 3128. This is the correct way to achieve the task. The other options are incorrect because they either delete a rule (-D), use the wrong protocol (top instead of tcp), or use the wrong port (81 instead of 80). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 381.

NEW QUESTION 164
A systems administrator wants to check for running containers. Which of the following commands can be used to show this information?

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

Answer: C

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

NEW QUESTION 165
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 170
A cloud engineer is asked to copy the file deployment.yaml from a container to the host where the container is running. Which of the following commands can accomplish this task?

A. docker cp container_id/deployment.yaml deployment.yaml
B. docker cp container_id:/deployment.yaml deployment.yaml
C. docker cp deployment.yaml local://deployment.yaml
D. docker cp container_id/deployment.yaml local://deployment.yaml

Answer: B

Explanation:
The command docker cp container_id:/deployment.yaml deployment.yaml can accomplish the task of copying the file deployment.yaml from a container to the host.
The docker command is a tool for managing Docker containers and images. The cp option copies files or directories between a container and the local filesystem. The container_id is the identifier of the container, which can be obtained by using the docker ps command.
The /deployment.yaml is the path of the file in the container, which must be preceded by a slash. The deployment.yaml is the path of the file on the host, which can be relative or absolute. The command docker cp container_id:/deployment.yaml deployment.yaml will copy the file deployment.yaml from the container to the current working directory on the host. This is the correct command to use to accomplish the task. The other options are incorrect because they either use the wrong syntax (docker cp container_id/deployment.yaml deployment.yaml or docker cp container_id/deployment.yaml local://deployment.yaml) or do not exist (docker cp deployment.yaml local://deployment.yaml). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 567.

NEW QUESTION 173
A Linux administrator needs to connect securely to a remote server in order to install application software. Which of the following commands would allow this connection?

A. scp "ABC-key.pem" root@10.0.0.1
B. sftp rooteiO.0.0.1
C. telnet 10.0.0.1 80
D. ssh -i "ABC-key.pem" root@10.0.0.1
E. sftp "ABC-key.pem" root@10.0.0.1

Answer: D

Explanation:
The command ssh -i "ABC-key.pem" root@10.0.0.1 would allow the administrator to connect securely to the remote server in order to install application software. The ssh command is a tool for establishing secure and encrypted connections between remote systems. The -i option specifies the identity file that contains the private key for key-based authentication. The "ABC-key.pem" is the name of the identity file that contains the private key. The root@10.0.0.1 is the username and the IP address of the remote server. The command ssh -i "ABC-key.pem" root@10.0.0.1 will connect to the remote server using the private key and allow the administrator to install application software. This is the correct command to use to connect securely to the remote server. The other options are incorrect because they either do not use key-based authentication (sftp root@10.0.0.1 or telnet 10.0.0.1 80) or do not use the correct syntax for the command (scp "ABC-key.pem" root@10.0.0.1 instead of scp -i "ABC-key.pem" root@10.0.0.1 or sftp "ABC-key.pem" root@10.0.0.1 instead of sftp -i "ABC-key.pem" root@10.0.0.1). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Implementing Basic Security, page 513.

NEW QUESTION 178
A systems administrator received a request to change a user's credentials. Which of the following commands will grant the request?

A. sudo passwd
B. sudo userde 1
C. sudo chage
D. sudo usermod

Answer: A

Explanation:
This command will allow the systems administrator to change the password of another user account in the system. The sudo prefix will grant the administrator the necessary privileges to perform this action, and the passwd command will prompt for the new password for the specified user. For example, if the administrator wants to change the password of a user named tom, the command will look like this:
sudo passwd tom
The other options are incorrect because:
* B. sudo userdel
This command will delete a user account from the system, not change its credentials. The userdel command removes the user's entry from the /etc/passwd and /etc/shadow files, as well as deletes the user's home directory and mail spool. This is not what the request asked for.
* C. sudo chage
This command will change the password expiration and aging information for a user account, not its credentials. The chage command can be used to set or modify various parameters related to password aging, such as the minimum and maximum number of days between password changes, the number of days before password expiration to issue a warning, and so on. This is not what the request asked for.

* D. sudo usermod
This command will modify various attributes of a user account, such as its login name, home directory, default shell, primary group, and so on. However, it cannot change the user's password directly. To do that, the usermod command requires the -p option followed by an encrypted password string, which is not easy to generate manually. Therefore, this is not a practical way to change a user's credentials.
References:
? How to Change Account Passwords on Linux
? How to Change a Password in Linux for Root and Other Users
? CompTIA Linux+ Certification Exam Objectives

NEW QUESTION 181
One leg of an LVM-mirrored volume failed due to the underlying physical volume, and a systems administrator is troubleshooting the issue. The following output has been provided:

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

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

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

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

Answer: B

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

NEW QUESTION 185
An administrator would like to list all current containers, regardless of their running state. Which of the following commands would allow the administrator to accomplish this task?

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

Answer: A

Explanation:
The best command to use to list all current containers, regardless of their running state, is A. docker ps -a. This command will show all containers, both running and stopped, with details such as container ID, image name, status, and ports. The other commands are either invalid or not relevant for this task. For example:
? B. docker list is not a valid command. There is no subcommand named list in docker.
? C. docker image ls will list all the images available on the local system, not the containers.
? D. docker inspect image will show detailed information about a specific image, not all the containers.

NEW QUESTION 188
An administrator accidentally deleted the /boot/vmlinuz file and must resolve the issue before the server is rebooted. Which of the following commands should the administrator use to identify the correct
version of this file?

A. rpm -qa | grep kernel; uname -a
B. yum -y update; shutdown -r now
C. cat /etc/centos-release; rpm -Uvh --nodeps
D. telinit 1; restorecon -Rv /boot

Answer: A

Explanation:
The command rpm -qa | grep kernel lists all the installed kernel packages, and the command uname -a displays the current kernel version. These commands can help the administrator identify the correct version of the /boot/vmlinuz file, which is the kernel image file. The other options are not relevant or helpful for this task. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing the Linux Boot Process, page 267.

NEW QUESTION 190
A Linux administrator is tasked with adding users to the system. However, the administrator wants to ensure the users' access will be disabled once the project is over. The expiration date should be 2021-09-30. Which of the following commands will accomplish this task?

A. sudo useradd -e 2021-09-30 Project_user
B. sudo useradd -c 2021-09-30 Project_user
C. sudo modinfo -F 2021-09-30 Project_uses
D. sudo useradd -m -d 2021-09-30 Project_user

Answer: A

Explanation:
The command that will accomplish this task is sudo useradd -e 2021-09-30 Project_user. This command will create a new user account named Project_user with an expiration date of 2021-09-30. The -e option of useradd specifies the date on which the user account will be disabled in YYYY-MM-DD format.
The other options are not correct commands for creating a user account with an expiration date. The sudo useradd -c 2021-09-30 Project_user command will create a new user account named Project_user with a comment of 2021-09-30. The -c option of useradd specifies a comment or description for the user account, not an expiration date. The sudo modinfo -F 2021-09-30 Project_user command is invalid because modinfo is not a command for managing user accounts, but a command for displaying information about kernel modules. The -F option of modinfo specifies a field name to show, not an expiration date. The sudo useradd -m -d 2021-09-30 Project_user command will create a new user account named Project_user with a home directory of 2021-09-30. The -m option of useradd specifies that the home directory should be created if it does not exist, and the -d option specifies the home directory name, not an expiration date. References: useradd(8) - Linux manual page; modinfo(8) - Linux manual page

NEW QUESTION 194
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 197
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 199
An administrator transferred a key for SSH authentication to a home directory on a remote server. The key file was moved to .ssh/authorized_keys location in order to establish SSH connection without a password. However, the SSH command still asked for the password. Given the following output:

```
[admin@linux ~ ]$ -ls -lhZ .ssh/auth*
-rw-r--r--. admin unconfined_u:object_r:user_home_t:s0 .ssh/authorized_keys
```

Which of the following commands would resolve the issue?

A. restorecon .ssh/authorized_keys
B. ssh_keygen -t rsa -o .ssh/authorized_keys
C. chown root:root .ssh/authorized_keys
D. chmod 600 .ssh/authorized_keys

Answer: D

Explanation:
The command that would resolve the issue is chmod 600 .ssh/authorized_keys. This command will change the permissions of the .ssh/authorized_keys file to 600, which means that only the owner of the file can read and write it. This is necessary for SSH key authentication to work properly, as SSH will refuse to use a key file that is accessible by other users or groups for security reasons. The output of ls -l shows that currently the .ssh/authorized_keys file has permissions of

664, which means that both the owner and group can read and write it, and others can read it.
The other options are not correct commands for resolving the issue. The restorecon .ssh/authorized_keys command will restore the default SELinux security context for the .ssh/authorized_keys file, but this will not change its permissions or ownership. The ssh_keygen -t rsa -o .ssh/authorized_keys command is invalid because ssh_keygen is not a valid command (the correct command is ssh-keygen), and the -o option is used to specify a new output format for the key file, not the output file name. The chown root:root
.ssh/authorized_keys command will change the owner and group of the .ssh/authorized_keys file to root, but this will not change its permissions or make it accessible by the user who wants to log in with SSH key authentication. References: How to Use Public Key Authentication with SSH; chmod(1) - Linux manual page

NEW QUESTION 203
A Linux administrator is troubleshooting an issue in which users are not able to access https://portal.comptia.org from a specific workstation. The administrator runs a few commands and receives the following output:

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

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

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

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

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

Answer: B

Explanation:
The best task to perform to resolve this issue is B. Remove the entry for portal.comptia.org from the local hosts file. This is because the local hosts file has a wrong entry that maps portal.comptia.org to 10.10.10.55, which is different from the actual IP address of 192.168.1.55 that is returned by the DNS server. This causes a mismatch and prevents the workstation from accessing the website. By removing or correcting the entry in the hosts file, the workstation will use the DNS server to resolve the domain name and access the website successfully.
To remove or edit the entry in the hosts file, you need to have root privileges and use a text editor such as vi or nano. For example, you can run the command:
sudo vi /etc/hosts
and delete or modify the line that says: 10.10.10.55 portal.comptia.org
Then save and exit the file.

NEW QUESTION 204
Which of the following files holds the system configuration for journal when running systemd?

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

Answer: A

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

NEW QUESTION 205
Users have reported that the interactive sessions were lost on a Linux server. A Linux administrator verifies the server was switched to rescue.target mode for maintenance. Which of the following commands will restore the server to its usual target?

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

Answer: B

Explanation:
The systemct1 reboot command will restore the server to its usual target by rebooting it. This will cause the server to load the default target specified in /etc/systemd/system.conf or /etc/systemd/system/default.target files. The telinit 0 command would shut down the server, not restore it to its usual target. The systemct1 get-default command would display the default target, not change it. The systemct1 emergency command would switch the server to emergency.target mode, which is even more

restrictive than rescue.target mode. References: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 17: System Maintenance and Operation, page 516.

NEW QUESTION 207
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 212
Employees in the finance department are having trouble accessing the file /opt/work/file. All IT employees can read and write the file. Systems administrator reviews the following output:

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

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

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

Answer: D

Explanation:
The command setfacl -m g:finance:rw file will permanently fix the access issue while limiting access to IT and finance department employees. The setfacl command is a tool for modifying the access control lists (ACLs) of files and directories on Linux systems. The ACLs are a mechanism that allows more fine-grained control over the permissions of files and directories than the traditional owner-group-others model. The - m option specifies the modification to the ACL. The g:finance:rw means that the group named finance will have read and write permissions on the file. The file is the name of the file to modify, in this case /opt/work/file. The command setfacl -m g:finance:rw file will add an entry to the ACL of the file that will grant read and write access to the finance group.
This will fix the access issue and allow the finance employees to access the file. The command will also preserve the existing permissions of the file, which means that the IT employees will still have read and write access to the file. This will limit the access to IT and finance department employees and prevent unauthorized access from other users.
This is the correct command to use to accomplish the task. The other options are incorrect because they either do not fix the access issue (chattr +i file or chown it:finance file) or do not limit the access to IT and finance department employees (chmod 666 file). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing File Permissions and Ownership, page 352.

NEW QUESTION 217
A systems administrator is tasked with creating a cloud-based server with a public IP address.

```
---
-name: start an instance with a public IP address
 community.abc.ec2_instance:
    name: "public-compute-instance"
    key_name: "comptia-ssh-key"
    vpc_subnet_id: subnet-5cjssh1
    instance_type: instance.type
    security_group: comptia
    network:
       assign_public_ip: true
    image_id: ami-1234568
    tags:
       Environment: Comptia-Items-Writing-Workshop
...
```

Which of the following technologies did the systems administrator use to complete this task?

A. Puppet
B. Git
C. Ansible
D. Terraform

Answer: D

Explanation:
The systems administrator used Terraform to create a cloud-based server with a public IP address. Terraform is a tool for building, changing, and versioning infrastructure as code. Terraform can create and manage resources on different cloud platforms, such as AWS, Azure, or Google Cloud. Terraform uses a declarative syntax to describe the desired state of the infrastructure and applies the changes accordingly. Terraform can also assign a public IP address to a cloud server by using the appropriate resource attributes. This is the correct technology that the systems administrator used to complete the task. The other options are incorrect because they are either not designed for creating cloud servers (Puppet or Git) or not capable of assigning public IP addresses (Ansible). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 559.

NEW QUESTION 218
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 223
.....

Relate Links

100% Pass Your XK0-005 Exam with Exambible Prep Materials

https://www.exambible.com/XK0-005-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/