CompTIA.XK0-005.v2024-03-22.q272

Exam Code:	XK0-005
Exam Name:	CompTIA Linux+ Certification Exam
Certification Provider:	CompTIA
Free Question Number:	272
Version:	v2024-03-22
# of views:	2140
# of Questions views:	2720
https://www.freeqas.com/qa/CompTIA/XK0-005/CompTIA.XK0-005.v2024-03-22.q272.html	

NEW QUESTION: 1
Joe, a user, is unable to log in to the Linux system Given the following output:

```
# grep joe /etc/passwd /etc/shadow
/etc/passwd:joe:x:1001:1001::/home/joe:/bin/nologin
/etc/shadow:joe:$6$3uOw6qWx9876jGhgKJsdfH987634534voj.:18883:0:99999:7:::
```

Which of the following command would resolve the issue?
A. usermod -s /bin/bash joe
B. pam_tally2 -u joe -r
C. passwd -u joe
D. chage -E 90 joe
Answer: (SHOW ANSWER)
Explanation
Based on the output of the image sent by the user, Joe is unable to log in to the Linux system because his account has been locked due to too many failed login attempts. The pam_tally2 -u joe -r command will resolve this issue by resetting Joe's failed login counter to zero and unlocking his account. This command uses the pam_tally2 module to manage user account locking based on login failures. The usermod -s /bin/bash joe command will change Joe's login shell to /bin/bash, but this will not unlock his account. The passwd -u joe command will unlock Joe's password if it has been locked by passwd -l joe, but this will not reset his failed login counter or unlock his account if it has been locked by pam_tally2. The chage -E 90 joe command will set Joe's account expiration date to 90 days from today, but this will not unlock his account or reset his failed login counter. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 537.

NEW QUESTION: 2
Which of the following files holds the system configuration for journal when running systemd?
A. /usr/lib/systemd/journalctl.conf
B. /etc/systemd/systemd-journald.conf
C. /etc/systemd/journald.conf
D. /etc/systemd/systemd-journalctl.conf

Answer: C (LEAVE A REPLY)

NEW QUESTION: 3
Which of the following technologies can be used as a central repository of Linux users and groups?
A. MFA
B. PAM
C. LDAP
D. SSO
Answer: (SHOW ANSWER)

NEW QUESTION: 4
An administrator created an initial Git repository and uploaded the first files. The administrator sees the following when listing the repository:

The administrator notices the file . DS STORE should not be included and deletes it from the online repository. Which of the following should the administrator run from the root of the local repository before the next commit to ensure the file is not uploaded again in future commits?
A. rm -f .DS STORE && git push
B. git fetch && git checkout .DS STORE
C. rm -f .DS STORE && git rebase origin main
D. echo .DS STORE >> .gitignore
Answer: (SHOW ANSWER)
Explanation
The correct answer is D. The administrator should run "echo .DS STORE >> .gitignore" from the root of the local repository before the next commit to ensure the file is not uploaded again in future commits.
This command will append the file name .DS STORE to the end of the .gitignore file, which is a special file that tells Git to ignore certain files or directories that should not be tracked or uploaded to the repository. By adding .DS STORE to the .gitignore file, the administrator will prevent Git from staging, committing, or pushing this file in the future.
The other options are incorrect because:
A: rm -f .DS STORE && git push
This command will delete the file .DS STORE from the local repository and then push the changes to the remote repository. However, this does not prevent the file from being uploaded again in future commits, if it is recreated or copied to the local repository.
B: git fetch && git checkout .DS STORE
This command will fetch the latest changes from the remote repository and then restore the file .DS STORE from the remote repository to the local repository. This is not what the administrator wants to do, as this will undo the deletion of the file from the online repository.

C: rm -f .DS STORE && git rebase origin main
This command will delete the file .DS STORE from the local repository and then rebase the local branch onto the main branch of the remote repository. This will rewrite the commit history of the local branch and may cause conflicts or errors. This is not what the administrator wants to do, as this is a risky and unnecessary operation.

NEW QUESTION: 5
A junior developer is unable to access an application server and receives the following output:

```
[root@server1 ~]# ssh dev2@172.16.25.126
dev2@172.16.25.126's password:
Permission denied, please try again.
dev2@172.16.25.126's password:
Permission denied, please try again.
dev2@172.16.25.126's password:
Account locked due to 4 failed logins
Account locked due to 5 failed logins
Last login: Mon Apr 22 21:21:06 2021 from 172.16.16.52
```

The systems administrator investigates the issue and receives the following output:

```
[root@server1 ~]# pam_tally2 --user=dev2
Login Failures Latest failure From
dev2 5 04/22/21 21:22:37 172.16.16.52
```

Which of the following commands will help unlock the account?
A. Pam_tally2 --user=dev2 --quiet
B. pam_ tally2 --user=dev2
C. pam_tally2 --user+dev2 --quiet
D. pam_tally2 --user=dev2 --reset
Answer: D (LEAVE A REPLY)
Explanation
To unlock an account that has been locked due to login failures, the administrator can use the command pam_tally2 --user=dev2 --reset (D). This will reset the failure counter for the user "dev2" and allow the user to log in again. The other commands will not unlock the account, but either display or increase the failure count. References:
[CompTIA Linux+ Study Guide], Chapter 4: Managing Users and Groups, Section: Locking Accounts with pam_tally2
[How to Lock and Unlock User Account in Linux]

NEW QUESTION: 6
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   9.5   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  6415   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: (SHOW ANSWER)

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: 7

A systems administrator installed a new software program on a Linux server. When the systems administrator tries to run the program, the following message appears on the screen.

```
Hardware virtualization support is not available on this system.
Either is not present or disabled in the system's BIOS
```

Which of the following commands will allow the systems administrator to check whether the system supports virtualization?

A. dmidecode -s system-version

B. lscpu

C. sysctl -a

D. cat /sys/device/system/cpu/possible

Answer: B (LEAVE A REPLY)

Explanation

The command that will allow the systems administrator to check whether the system supports virtualization is lscpu. This command will display information about the CPU architecture, such as the number of CPUs, cores, sockets, threads, model name, frequency, cache size, and flags. One of the flags is vmx (for Intel processors) or svm (for AMD processors), which indicates that the CPU supports hardware virtualization. If the flag is present, it means that the system supports virtualization. If the flag is absent, it means that the system does not support virtualization or that it is disabled in the BIOS settings.

The other options are not correct commands for checking whether the system supports virtualization. The dmidecode -s system-version command will display the version of the system, such as the product name or serial number, but not the CPU information. The sysctl -a command will display all the kernel parameters, but not the CPU flags. The cat /sys/devices/system/cpu/possible command will display the range of possible CPUs that can be online or offline, but not the CPU features. References: lscpu(1) - Linux manual page; How To Check If Virtualization is Enabled in Windows 10 / 11

NEW QUESTION: 8
A user reported issues when trying to log in to a Linux server. The following outputs were received:
Given the outputs above. which of the following is the reason the user is una-ble to log in to the server?
A. User1 needs to set a long password.
B. User1 is in the incorrect group.
C. The user1 shell assignment incorrect.
D. The user1 password is expired.
Answer: D (LEAVE A REPLY)
Explanation
The user1 password is expired. This can be inferred from the output of the chage -l user1 command, which shows the password expiration information for user1. The output shows that the password expired on
2020-10-01, and the account expired on 2020-10-08. This means that user1 cannot log in to the server unless the password and account are reactivated by the system administrator.
The other options are not correct based on the outputs above. User1 does not need to set a long password, because the output of the passwd -S user1 command shows that the password has a minimum length of 5 characters, which is met by user1's password. User1 is not in the incorrect group, because the output of the groups user1 command shows that user1 belongs to the app group, which is presumably the correct group for accessing the server. The user1 shell assignment is not incorrect, because the output of the grep user1 /etc/passwd command shows that user1 has /bin/bash as the default shell, which is a valid and common shell for Linux users.

NEW QUESTION: 9
A Linux administrator cloned an existing Linux server and built a new server from that clone. The administrator encountered the following error after booting the cloned server:

```
Device mismatch detected
```

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

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


Which of the following should administrator use to resolve the device mismatch issue and mount the disk?
A. mount disk by device-id
B. fsck -A
C. mount disk by-label
D. mount disk by-blkid
Answer: (SHOW ANSWER)
Explanation

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

NEW QUESTION: 10
A Linux engineer set up two local DNS servers (10.10.10.10 and 10.10.10.20) and was testing email connectivity to the local mail server using the mail command on a local machine when the following error appeared:

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

The local machine DNS settings are:

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

Mail server: mail.example.com
```


Which of the following commands could the engineer use to query the DNS server to get mail server information?
A. dig @example.com 10.10.10.20 a
B. dig @10.10.10.20 example.com mx
C. dig @example.com 10.10.10.20 ptr
D. dig @10.10.10.20 example.com ns
Answer: (SHOW ANSWER)
Explanation
The command dig @10.10.10.20 example.com mx will query the DNS server to get mail server information.
The dig command is a tool for querying DNS servers and displaying the results. The @ option specifies the DNS server to query, in this case 10.10.10.20. The mx option specifies the type of record to query, in this case mail exchange (MX) records, which identify the mail servers for a domain. The domain name to query is example.com. This command will show the MX records for example.com from the DNS server 10.10.10.20.
This is the correct command to use to accomplish the task. The other options are incorrect because they either use the wrong syntax (@example.com 10.10.10.20 instead of @10.10.10.20 example.com), the wrong type of record (a or ptr instead of mx), or the wrong domain name (example.com ns instead of example.com mx). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 415.

NEW QUESTION: 11

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

A. hostnamectl status --no-ask-password

B. hostnamectl set-hostname "$(perl -le "print" "A" x 86)"

C. hostnamectl set-hostname Comptia-WebNode --transient

D. hostnamectl set-hostname Comptia-WebNode -H root@192.168.2.14

Answer: D (LEAVE A REPLY)

NEW QUESTION: 12

A Linux engineer needs to create a custom script, cleanup.sh, to run at boot as part of the system services. Which of the following processes would accomplish this task?

A. Create a unit file in the /etc/default/ directory.

systemctl enable cleanup

systemctl is-enabled cleanup

B. Create a unit file in the /etc/ske1/ directory.

systemctl enable cleanup

systemctl is-enabled cleanup

C. Create a unit file in the /etc/systemd/system/ directory.

systemctl enable cleanup

systemctl is-enabled cleanup

D. Create a unit file in the /etc/sysctl.d/ directory.

systemctl enable cleanup

systemctl is-enabled cleanup

Answer: (SHOW ANSWER)

The process that will accomplish the task of creating a custom script to run at boot as part of the system services is:

Create a unit file in the /etc/systemd/system/ directory. A unit file is a configuration file that defines the properties and behavior of a systemd service. The systemd is a system and service manager that controls the startup and operation of Linux systems.

The /etc/systemd/system/ directory is the location where the administrator can create and store custom unit files. The unit file should have a name that matches the name of the script, such as cleanup.service, and should contain the following sections and options:

[Unit]: This section provides the general information about the service, such as the description, dependencies, and conditions. The administrator should specify the following options in this section:

Description: A brief description of the service, such as "Custom cleanup script".

After: The name of another unit that this service should start after, such as "network.target".

ConditionPathExists: The path of the file or directory that must exist for the service to start, such as "/opt/scripts/cleanup.sh".

[Service]: This section defines how the service should be started and stopped, and what commands should be executed. The administrator should specify the following options in this section:

Type: The type of the service, such as "oneshot", which means that the service will run once and then exit.

ExecStart: The command that will start the service, such as "/bin/bash /opt/scripts/cleanup.sh".

RemainAfterExit: A boolean value that indicates whether the service should remain active after the command exits, such as "yes".

[Install]: This section defines how the service should be enabled and under what circumstances it should be started. The administrator should specify the following option in this section:

WantedBy: The name of another unit that wants this service to be started, such as "multi-user.target", which means that the service will be started when the system reaches the multi-user mode.
Run the command systemctl enable cleanup. This command will enable the service and create the necessary symbolic links to start the service at boot.
Run the command systemctl is-enabled cleanup. This command will check the status of the service and confirm that it is enabled.
This process will create a custom script, cleanup.sh, to run at boot as part of the system services. This is the correct process to use to accomplish the task. The other options are incorrect because they either use the wrong directory for the unit file (/etc/default/, /etc/skel/, or /etc/sysctl.d/) or do not create a unit file at all. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing System Services, pages 457-459.

NEW QUESTION: 13
Which of the following is the best tool for dynamic tuning of kernel parameters?
A. tuned
B. tune2fs
C. tuned-adm
D. turbostat
Answer: (SHOW ANSWER)
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: 14
A Linux administrator wants to find out whether files from the wget package have been altered since they were installed. Which of the following commands will provide the correct information?
A. rpm -i wget
B. rpm -qf wget
C. rpm -F wget
D. rpm -V wget
Answer: D (LEAVE A REPLY)
Explanation
The command that will provide the correct information about whether files from the wget package have been altered since they were installed is rpm -V wget. This command will use the rpm utility to verify an installed RPM package by comparing information about the installed files with information from the RPM database.
The verification process can check various attributes of each file, such as size, mode, owner, group, checksum, capabilities, and so on. If any discrepancies are found, rpm will report them using a single letter code for each attribute.

The other options are not correct commands for verifying an installed RPM package. The rpm -i wget command is invalid because -i is used to install a package from a file, not to verify an installed package. The rpm -qf wget command will query which package owns wget as a file name or path name, but it will not verify its attributes. The rpm -F wget command will freshen (upgrade) an already installed package with wget as a file name or path name, but it will not verify its attributes. References: rpm(8) - Linux manual page; Using RPM to Verify Installed Packages

NEW QUESTION: 15
A systems administrator created a web server for the company and is required to add a tag for the API so end users can connect. Which of the following would the administrator do to complete this requirement?
A. hostnamectl set-hostname "$(perl -le "print" "A" x 86)"
B. hostnamectl set-hostname Comptia-WebNode -H root@192.168.2.14
C. hostnamectl status --no-ask-password
D. hostnamectl set-hostname Comptia-WebNode --transient
Answer: B (LEAVE A REPLY)
Explanation
The command hostnamectl set-hostname Comptia-WebNode -H root@192.168.2.14 sets the hostname of the web server to Comptia-WebNode and connects to the server using the SSH protocol and the root user. This is the correct way to complete the requirement. The other options are incorrect because they either display the current hostname status (hostnamectl status), set an invalid hostname (hostnamectl set-hostname "$(perl -le
"print" "A" x 86)"), or set a transient hostname that is not persistent (hostnamectl set-hostname Comptia-WebNode --transient).
References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter
9: Managing System Components, page 291.

NEW QUESTION: 16
A Linux systems administrator needs to copy files and directories from Server A to Server B.
Which of the following commands can be used for this purpose? (Select TWO)
A. ssh
B. rsyslog
C. cp
D. rsync
E. reposync
F. scp
Answer: D,F (LEAVE A REPLY)

NEW QUESTION: 17
An administrator recently updated the BIND software package and would like to review the default configuration that shipped with this version. Which of the following files should the administrator review?
A. /etc/named.conf.rpmsave
B. /etc/named.conf
C. /etc/named.conf.rpmnew
D. /etc/bind/bind.conf
Answer: (SHOW ANSWER)
Explanation
After installing a new version of a package that includes a configuration file that already exists on the system, such as /etc/httpd/conf/httpd.conf, RPM will create a new file with the .rpmnew extension instead of overwriting the existing file. This allows the administrator to review the default configuration that shipped with this version and compare it with the current configuration before deciding whether to merge or replace the files. The /etc/named.conf.rpmsave file is created by RPM when a package is uninstalled and it contains a configuration file that was modified by the administrator. This allows the administrator to restore the configuration file if needed.
The /etc/named.conf file is the main configuration file for the BIND name server, not the httpd web server. The /etc/bind/bind.conf file does not exist by default in Linux systems. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Packages and Software, page 561.

NEW QUESTION: 18
A systems administrator is enabling LUKS on a USB storage device with an ext4 filesystem format. The administrator runs dmesg and notices the following output:

```
sd 8:0:0:0: [sdc] Attached SCSI disk
EXT4-fs (sdc1): mounting ext3 file system using the ext4 subsystem
EXT4-fs (sdc1): mounted filesystem with ordered data mode.  Opts: (null)
```

Given this scenario, which of the following should the administrator perform to meet these requirements?
(Select three).
A. gpg /dev/sdcl
B. pvcreate /dev/sdc
C. mkfs . ext4 /dev/mapper/LUKSCJ001 - L ENCRYPTED
D. umount / dev/ sdc
E. fdisk /dev/sdc
F. mkfs . vfat /dev/mapper/LUKS0001 - L ENCRYPTED
G. wipefs -a/dev/sdbl
H. cryptsetup luksFormat /dev/ sdcl
Answer: C,D,H (LEAVE A REPLY)
Explanation
To enable LUKS on a USB storage device with an ext4 filesystem format, the administrator needs to perform the following steps:
Unmount the device if it is mounted using umount /dev/sdc (D)
Create a partition table on the device using fdisk /dev/sdc (E)

Format the partition with LUKS encryption using cryptsetup luksFormat /dev/sdc1 (H) Open the encrypted partition using cryptsetup luksOpen /dev/sdc1 LUKS0001 Create an ext4 filesystem on the encrypted partition using mkfs.ext4 /dev/mapper/LUKS0001 Mount the encrypted partition using mount /dev/mapper/LUKS0001 /mnt References:
[CompTIA Linux+ Study Guide], Chapter 9: Securing Linux, Section: Encrypting Disks
[How to Encrypt USB Drive on Ubuntu 18.04]

NEW QUESTION: 19
A Linux administrator needs to create a new user named user02. However, user02 must be in a different home directory, which is under /comptia/projects. Which of the following commands will accomplish this task?
A. useradd -d /comptia/projects user02
B. useradd -s /comptia/projects user02
C. useradd -b /comptia/projects user02
D. useradd -m /comptia/projects user02
Answer: (SHOW ANSWER)

NEW QUESTION: 20
A systems administrator installed a new software program on a Linux server. When the systems administrator tries to run the program, the following message appears on the screen.

```
Hardware virtualization support is not available on this system.
Either is not present or disabled in the system's BIOS
```

Which of the following commands will allow the systems administrator to check whether the system supports virtualization?
A. dmidecode -s system-version
B. lscpu
C. sysctl -a
D. cat /sys/device/system/cpu/possible
Answer: B (LEAVE A REPLY)
Explanation
The command that will allow the systems administrator to check whether the system supports virtualization is lscpu. This command will display information about the CPU architecture, such as the number of CPUs, cores, sockets, threads, model name, frequency, cache size, and flags. One of the flags is vmx (for Intel processors) or svm (for AMD processors), which indicates that the CPU supports hardware virtualization. If the flag is present, it means that the system supports virtualization. If the flag is absent, it means that the system does not support virtualization or that it is disabled in the BIOS settings.
The other options are not correct commands for checking whether the system supports virtualization. The dmidecode -s system-version command will display the version of the system, such as the product name or serial number, but not the CPU information. The sysctl -a command will display all the kernel parameters, but not the CPU flags. The cat /sys/devices/system/cpu/possible command will display the range of possible CPUs that can be online or offline, but not the CPU features. References: lscpu(1) - Linux manual page; How To Check If Virtualization is Enabled in Windows 10 / 11

NEW QUESTION: 21

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 (LEAVE A REPLY)

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: 22

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. chmod -R 400 /home/web

B. setfacl read /home/web

C. chown web:web /home/web

D. echo "umask 377" >> /home/web/.bashrc

Answer: A (LEAVE A REPLY)

NEW QUESTION: 23

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. Whileworking 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 (LEAVE A REPLY)

Explanation

The command restorecon -R -v /var/www/html will best resolve the issue. The issue is caused by the incorrect SELinux context of the web server files under the /var/www/html directory. The output of ls -Z

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

NEW QUESTION: 24
A Linux system is failing to boot with the following error:

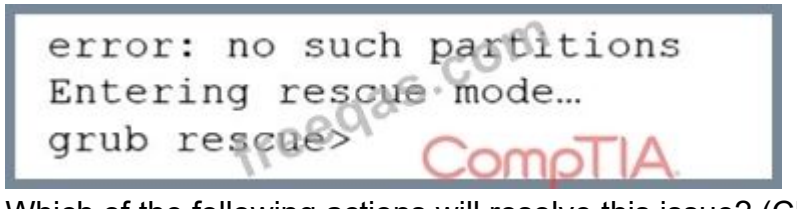

Which of the following actions will resolve this issue? (Choose two.)
A. Execute grub-install --root-directory=/mnt and reboot.
B. Execute grub-install /dev/sdX and reboot.
C. Interrupt the boot process in the GRUB menu and add rescue to the kernel line.
D. Fix the partition modifying /etc/default/grub and reboot.
E. Interrupt the boot process in the GRUB menu and add single to the kernel line.
F. Boot the system on a LiveCD/ISO.
Answer: B,F (LEAVE A REPLY)
Explanation
The administrator should do the following two actions to resolve the issue:
Boot the system on a LiveCD/ISO. This is necessary to access the system and repair the boot loader. A LiveCD/ISO is a bootable media that contains a Linux distribution that can run without installation. The administrator can boot the system from the LiveCD/ISO and mount the root partition of the system to a temporary directory, such as /mnt.
Execute grub-install /dev/sdX and reboot. This will reinstall the GRUB boot loader to the disk device, where sdX is the device name of the disk, such as sda or sdb. The GRUB boot loader is a program that runs when the system is powered on and allows the user to choose which operating system or kernel to boot. The issue is caused by a corrupted or missing GRUB boot loader, which prevents the system from booting. The command grub-install will restore the GRUB boot loader and fix the issue.
The other options are incorrect because they either do not fix the boot loader (interrupt the boot process in the GRUB menu or fix the partition modifying /etc/default/grub) or do not use the correct syntax (grub-install
--root-directory=/mnt instead of grub-install /dev/sdX or rescue or single instead of recovery in the GRUB menu). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing the Linux Boot Process, pages 265-266.

NEW QUESTION: 25
An application developer received a file with the following content:
##This is a sample Image ##
FROM ubuntu:18.04

MAINTAINER demohut@gtmail.com.hac
COPY . /app
RUN make /app
CMD python /app/app.py
RUN apt-get update
RUN apt-get install -y nginx
CMD ["echo","Image created"]
The developer must use this information to create a test bed environment and identify the image (myimage) as the first version for testing a new application before moving it to production. Which of the following commands will accomplish this task?
A. docker build -t myimage:1.0 .
B. docker build -t myimage: .
C. docker build -t myimage-1.0 .
D. docker build -i myimage:1.0 .
Answer: A (LEAVE A REPLY)
Explanation
The docker build command is used to build an image from a Dockerfile and a context1. The Dockerfile is a text file that contains the instructions for creating the image, and the context is a set of files that can be used in the image creation process1. The file that the developer received is an example of a Dockerfile.
The -t option is used to specify a name and an optional tag for the image1. The name and tag are separated by a colon (:), and the tag is usually used to indicate the version of the image2. For example, -t myimage:1.0 means that the image will be named myimage and tagged as 1.0.
The last argument of the docker build command is the path to the context, which can be a local directory or a URL1. The dot (.) means that the current working directory is the context2. Therefore, docker build -t myimage:1.0 . means that the image will be built from the Dockerfile and the files in the current working directory, and it will be named myimage and tagged as 1.0.

NEW QUESTION: 26
A systems administrator is enabling LUKS on a USB storage device with an ext4 filesystem format. The administrator runs dmesg and notices the following output:

```
sd 8:0:0:0: [sdc] Attached SCSI disk
EXT4-fs (sdc1): mounting ext3 file system using the ext4 subsystem
EXT4-fs (sdc1): mounted filesystem with ordered data mode.  Opts: (null)
```

Given this scenario, which of the following should the administrator perform to meet these requirements?
(Select three).
A. gpg /dev/sdcl
B. pvcreate /dev/sdc
C. mkfs . ext4 /dev/mapper/LUKSCJ001 - L ENCRYPTED
D. umount / dev/ sdc
E. fdisk /dev/sdc
F. mkfs . vfat /dev/mapper/LUKS0001 - L ENCRYPTED
G. wipefs -a/dev/sdbl

H. cryptsetup luksFormat /dev/ sdcl
Answer: (SHOW ANSWER)
Explanation
To enable LUKS on a USB storage device with an ext4 filesystem format, the administrator needs to perform the following steps:
Unmount the device if it is mounted using umount /dev/sdc (D)
Create a partition table on the device using fdisk /dev/sdc (E)
Format the partition with LUKS encryption using cryptsetup luksFormat /dev/sdc1 (H) Open the encrypted partition using cryptsetup luksOpen /dev/sdc1 LUKS0001 Create an ext4 filesystem on the encrypted partition using mkfs.ext4 /dev/mapper/LUKS0001 Mount the encrypted partition using mount /dev/mapper/LUKS0001 /mnt References:
[CompTIA Linux+ Study Guide], Chapter 9: Securing Linux, Section: Encrypting Disks
[How to Encrypt USB Drive on Ubuntu 18.04]

NEW QUESTION: 27
Which of the following tools is BEST suited to orchestrate a large number of containers across many different servers?
A. Kubernetes
B. Ansible
C. Podman
D. Terraform
Answer: (SHOW ANSWER)
The tool that is best suited to orchestrate a large number of containers across many different servers is Kubernetes. Kubernetes is an open-source platform for managing containerized applications and services. Kubernetes allows the administrator to deploy, scale, and update containers across a cluster of servers, as well as to automate the configuration and coordination of the containers. Kubernetes also provides features such as service discovery, load balancing, storage management, security, monitoring, and logging. Kubernetes can handle complex and dynamic workloads and ensure high availability and performance of the containers. Kubernetes is the tool that is best suited to orchestrate a large number of containers across many different servers. This is the correct answer to the question. The other options are incorrect because they either do not orchestrate containers (Ansible or Terraform) or do not operate across many different servers (Podman). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 573.

NEW QUESTION: 28
A Linux administrator is troubleshooting an issue in which an application service failed to start on a Linux server. The administrator runs a few commands and gets the following outputs:

Output 1:

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

Output 2:

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


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

A. Enable the logsearch.service and restart the service.

B. Increase the TimeoutStartUSec configuration for the logsearch.sevice.

C. Update the OnCalendar configuration to schedule the start of the logsearch.service.

D. Update the KillSignal configuration for the logsearch.service to use TERM.

Answer: B (LEAVE A REPLY)

Explanation

The administrator should increase the TimeoutStartUSec configuration for the logsearch.service to resolve the issue. The output of systemct1 status logsearch.service shows that the service failed to start due to a timeout.

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

NEW QUESTION: 29

A Linux engineer needs to create a custom script, cleanup.sh, to run at boot as part of the system services.

Which of the following processes would accomplish this task?

A. Create a unit file in the /etc/default/ directory.

systemct1 enable cleanup

systemct1 is-enabled cleanup

B. Create a unit file in the /etc/ske1/ directory.

systemct1 enable cleanup

systemct1 is-enabled cleanup

C. Create a unit file in the /etc/systemd/system/ directory.

systemct1 enable cleanup

systemct1 is-enabled cleanup

D. Create a unit file in the /etc/sysctl.d/ directory.

systemct1 enable cleanup

systemct1 is-enabled cleanup

Answer: C (LEAVE A REPLY)

Explanation

The process that will accomplish the task of creating a custom script to run at boot as part of the system services is:

Create a unit file in the /etc/systemd/system/ directory. A unit file is a configuration file that defines the properties and behavior of a systemd service. The systemd is a system and service manager that controls the startup and operation of Linux systems.

The /etc/systemd/system/ directory is the location where the administrator can create and store custom unit files. The unit file should have a name that matches the name of the script, such as cleanup.service, and should contain the following sections and options:

[Unit]: This section provides the general information about the service, such as the description, dependencies, and conditions. The administrator should specify the following options in this section:

Description: A brief description of the service, such as "Custom cleanup script".

After: The name of another unit that this service should start after, such as "network.target".

ConditionPathExists: The path of the file or directory that must exist for the service to start, such as "/opt/scripts/cleanup.sh".

[Service]: This section defines how the service should be started and stopped, and what commands should be executed. The administrator should specify the following options in this section:

Type: The type of the service, such as "oneshot", which means that the service will run once and then exit.

ExecStart: The command that will start the service, such as "/bin/bash /opt/scripts/cleanup.sh".

RemainAfterExit: A boolean value that indicates whether the service should remain active after the command exits, such as "yes".

[Install]: This section defines how the service should be enabled and under what circumstances it should be started. The administrator should specify the following option in this section:

WantedBy: The name of another unit that wants this service to be started, such as "multi-user.target", which means that the service will be started when the system reaches the multi-user mode.

Run the command systemct1 enable cleanup. This command will enable the service and create the necessary symbolic links to start the service at boot.

Run the command systemct1 is-enabled cleanup. This command will check the status of the service and confirm that it is enabled.

This process will create a custom script, cleanup.sh, to run at boot as part of the system services. This is the correct process to use to accomplish the task. The other options are incorrect because they either use the wrong directory for the unit file (/etc/default/, /etc/skel/, or /etc/sysctl.d/) or do not create a unit file at all. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing System Services, pages 457-459.

NEW QUESTION: 30

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

```
Oct 20 03:45:50 hostname kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx 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 (LEAVE A REPLY)

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: 31

A developer reported an incident involving the application configuration file /etc/httpd/conf/httpd.conf that is missing from the server. Which of the following identifies the RPM package that installed the configuration file?

A. rpm -qf /etc/httpd/conf/httpd.conf

B. rpm -ql /etc/httpd/conf/httpd.conf

C. rpm -query /etc/httpd/conf/httpd.conf

D. rpm -q /etc/httpd/conf/httpd.conf

Answer: A (LEAVE A REPLY)

The rpm -qf /etc/httpd/conf/httpd.conf command will identify the RPM package that installed the configuration file. This command will query the database of installed packages and display the name of the package that owns the specified file. The rpm -ql /etc/httpd/conf/httpd.conf command is invalid, as -ql is not a valid option for rpm. The rpm --query /etc/httpd/conf/httpd.conf command is incorrect, as --query requires a package name, not a file name. The rpm -q /etc/httpd/conf/httpd.conf command is incorrect, as -q requires a package name, not a file name. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Packages and Software, page 560.

NEW QUESTION: 32

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

A. ls

B. du

C. dd

D. df

E. parted

F. mount

G. fdisk

Answer: (SHOW ANSWER)

NEW QUESTION: 33

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

A. ip link show eth1

B. netstat -r eth1

C. ss -ti eth1

D. ifconfig hw eth1

Answer: A (LEAVE A REPLY)

NEW QUESTION: 34
A Linux user reported the following error after trying to connect to the system remotely:
ssh: connect to host 10.0.1.10 port 22: Resource temporarily unavailable The Linux systems administrator executed the following commands in the Linux system while trying to diagnose this issue:

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

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


Which of the following commands will resolve this issue?
A. firewall-cmd --zone=public --permanent --add-service=ssh
B. firewall-cmd --zone=public --permanent --add-service=22
C. systemctl enable firewalld; systemctl restart firewalld
D. firewall-cmd --zone=public --permanent --add-port=22/udp
Answer: A (LEAVE A REPLY)

NEW QUESTION: 35
Which of the following would significantly help to reduce data loss if more than one drive fails at the same time?
A. Server clustering
B. Load balancing
C. RAID
D. VDI
Answer: C (LEAVE A REPLY)
Explanation
RAID stands for Redundant Array of Independent Disks, which is a technology that combines multiple physical disks into a logical unit that provides improved performance, reliability, or both. RAID can significantly help to reduce data loss if more than one drive fails at the same time, depending on the RAID level used. For example, RAID 1 (mirroring) duplicates the data on two or more disks, so that if one disk fails, the data can be recovered from another disk. RAID 5 (striping with parity) distributes the data and parity information across three or more disks, so that if one disk fails, the data can be reconstructed from the remaining disks. RAID 6 (striping with double parity) extends RAID 5

by adding another parity block, so that if two disks fail, the data can still be recovered from the remaining disks. References: [What is RAID?]

NEW QUESTION: 36
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. sudo vi check
C. cat /etc/sudoers | tee test
D. test -f /etc/sudoers
Answer: A (LEAVE A REPLY)
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). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 546.

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

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

Which of the following commands would resolve the issue and allow an SSH connection to be established without a password?
A. restorecon -rv .ssh/authorized_key
B. mv .ssh/authorized_key .ssh/authorized_keys
C. systemct1 restart sshd.service
D. chmod 600 mv .ssh/authorized_key
Answer: B (LEAVE A REPLY)
Explanation
The command mv .ssh/authorized_key .ssh/authorized_keys will resolve the issue and allow an SSH connection to be established without a password. The issue is caused by the incorrect file name of the authorized key file on the server. The file should be named authorized_keys, not authorized_key.
The mv command will rename the file and fix the issue. The other options are incorrect because they either do not affect the file name (restorecon or chmod) or do not restart the SSH service (systemct1). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 410.

NEW QUESTION: 38
Joe, a user, is unable to log in to the Linux system Given the following output:

```
# grep joe /etc/passwd /etc/shadow
/etc/passwd:joe:x:1001:1001::/home/joe:/bin/nologin
/etc/shadow:joe:$6$3uOw6qWx9876jGhgKJsdfH987634534voj.:18883:0:99999:7:::
```

Which of the following command would resolve the issue?

A. usermod -s /bin/bash joe

B. pam_tally2 -u joe -r

C. passwd -u joe

D. chage -E 90 joe

Answer: B (LEAVE A REPLY)

Based on the output of the image sent by the user, Joe is unable to log in to the Linux system because his account has been locked due to too many failed login attempts. The pam_tally2 -u joe -r command will resolve this issue by resetting Joe's failed login counter to zero and unlocking his account. This command uses the pam_tally2 module to manage user account locking based on login failures. The usermod -s /bin/bash joe command will change Joe's login shell to /bin/bash, but this will not unlock his account. The passwd -u joe command will unlock Joe's password if it has been locked by passwd -l joe, but this will not reset his failed login counter or unlock his account if it has been locked by pam_tally2. The chage -E 90 joe command will set Joe's account expiration date to 90 days from today, but this will not unlock his account or reset his failed login counter. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 537.

NEW QUESTION: 39

Which of the following data structures is written in JSON?

A)

```
---
name: user1
position: DevOps
floor: 3
```

B)

```
<table>
<tbody><tr>
<td>user1</td>
<td>DevOps</td>
<td>3</td>
</tr>
</tbody></table>
```

C)

```
<root>
   <floor>3</floor>
   <name>user1</name>
   <position>DevOps</position>
</root>
```

D)

```
{
   "name": "user1",
   "job": "DevOps",
   "floor": 3
}
```

A. Option A
B. Option B
C. Option D
D. Option C
Answer: C (LEAVE A REPLY)

NEW QUESTION: 40
A Linux administrator rebooted a server. Users then reported some of their files were missing. After doing some troubleshooting, the administrator found one of the filesystems was missing. The filesystem was not listed in /etc/f stab and might have been mounted manually by someone prior to reboot. Which of the following would prevent this issue from reoccurring in the future?
A. Mount the filesystem manually.
B. Sync the mount units.
C. Create a mount unit and enable it to be started at boot.
D. Remount all the missing filesystems
Answer: (SHOW ANSWER)

NEW QUESTION: 41
An administrator attempts to rename a file on a server but receives the following error.

```
mv: cannot move 'files/readme.txt' to 'files/readme.txt.orig': Operation not permitted.
```

The administrator then runs a few commands and obtains the following output:

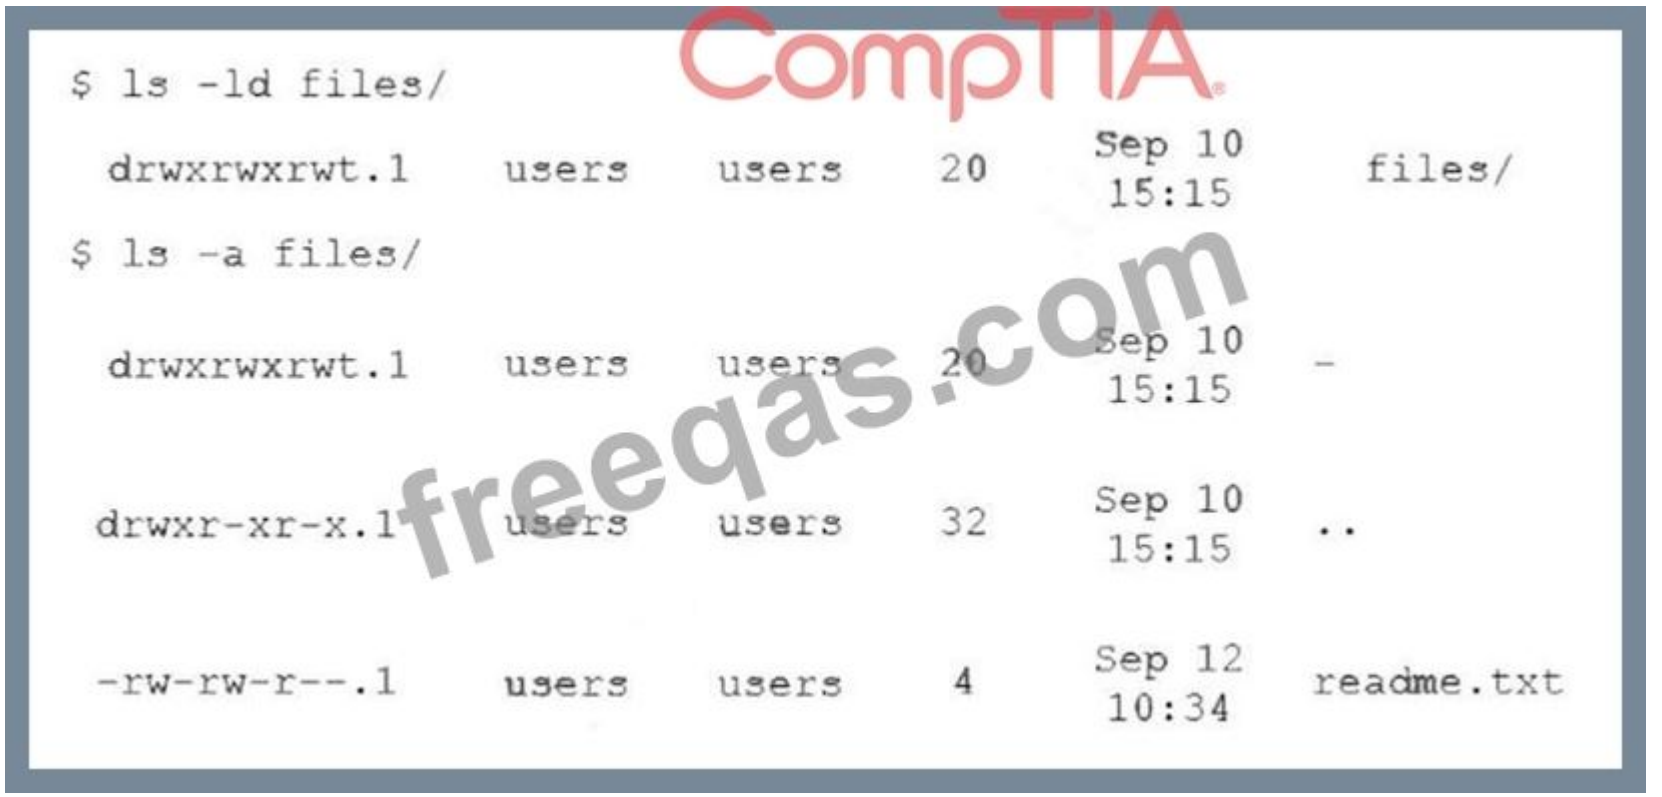

Which of the following commands should the administrator run NEXT to allow the file to be renamed by any user?

A. chgrp reet files

B. chacl -R 644 files

C. chown users files

D. chmod -t files

Answer: D (LEAVE A REPLY)

Explanation

The command that the administrator should run NEXT to allow the file to be renamed by any user is chmod -t files. This command uses the chmod tool, which is used to change file permissions and access modes. The -t option removes (or sets) the sticky bit on a directory, which restricts deletion or renaming of files within that directory to only their owners or root. In this case, since files is a directory with sticky bit set (indicated by t in drwxrwxrwt), removing it will allow any user to rename or delete files within that directory.

The other options are not correct commands for allowing any user to rename files within files directory. The chgrp reet files command will change the group ownership of files directory to reet, but it will not affect its permissions or access modes. The chacl -R 644 files command is invalid, as chacl is used to change file access control lists (ACLs), not permissions or access modes. The chown users files command will change the user ownership of files directory to users, but it will not affect its permissions or access modes. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing Users and Groups; chmod(1) - Linux manual page

NEW QUESTION: 42

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 -j INPUT 192.168.10.50 -p DROP

B. iptables -i INPUT --ipv4 192.168.10.50 -z DROP

C. iptables -F INPUT -j 192.168.10.50 -m DROP

D. iptables -A INPUT -s 192.168.10.30 -j DROP

Answer: D (LEAVE A REPLY)

NEW QUESTION: 43
A systems administrator is implementing a new service task with systems at startup and needs to execute a script entitled test.sh with the following content:

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

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

NEW QUESTION: 44
A systems administrator is notified that the mysqld process stopped unexpectedly. The systems administrator issues the following command:
sudo grep -i -r 'out of memory' /var/log
The output of the command shows the following:
kernel: Out of memory: Kill process 9112 (mysqld) score 511 or sacrifice child.

Which of the following commands should the systems administrator execute NEXT to troubleshoot this issue? (Select two).

A. free -h

B. nc -v 127.0.0.1 3306

C. renice -15 $(pidof mysql)

D. lsblk

E. killall -15

F. vmstat -a 1 4

Answer: (SHOW ANSWER)

The free -h command can be used to check the amount of free and used memory in the system in a human-readable format. This can help to troubleshoot the issue of mysqld being killed due to out of memory. The vmstat -a 1 4 command can be used to monitor the system's virtual memory statistics, such as swap usage, paging activity, and memory faults, every one second for four times. This can help to identify any memory pressure or performance issues that may cause out of memory errors. The nc -v 127.0.0.1 3306 command would attempt to connect to the MySQL server on port 3306 and display any diagnostic messages, but this would not help to troubleshoot the memory issue. The renice -15 $(pidof mysql) command would change the priority of the mysql process to -15, but this would not prevent it from being killed due to out of memory. The lsblk command would display information about block devices, not memory usage. The killall -15 command would send a SIGTERM signal to all processes with a matching name, but this would not help to troubleshoot the memory issue. Reference: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 15: Managing Memory and Process Execution, pages 468-469.

NEW QUESTION: 45

When trying to log in remotely to a server, a user receives the following message:

```
Password:
Last failed login: Wed Sep 15 17:23:45 CEST 2021 from 10.0.4.3 on ssh:notty
There were 3 failed login attempts since the last successful login.
Connection to localhost closed.
```

The server administrator is investigating the issue on the server and receives the following outputs:

Output 1:

```
user:x:1001:1004::/home/user:/bin/false
```

Output 2:

```
dzwx------. 2 user 62 Sep 15 17:17 /home/user
```

Output 3:

```
Sep 12 14:14:05 server sshd[22958] Failed password for user from 10.0.2.8
Sep 15 17:24:03 server sshd[8460]: Accepted keyboard-interactive/pam for user from 10.0.6.5 port 50928 ssh2
Sep 15 17:24:03 server sshd[8460]: pam_unix(sshd:session): session opened for user testuser
Sep 15 17:24:03 server sshd[8460]: pam_unix(sshd:session): session closed for user testuser
```

Which of the following is causing the issue?

A. The wrong permissions are on the user's home directory.

B. The account was locked out due to three failed logins.

C. The user entered the wrong password.
D. The user has the wrong shell assigned to the account.
Answer: (SHOW ANSWER)
Explanation
The user has the wrong shell assigned to the account, which is causing the issue. The output 1 shows that the user's shell is set to /bin/false, which is not a valid shell and will prevent the user from logging in. The output 2 shows that the user's home directory has the correct permissions (drwxr-xr-x), and the output 3 shows that the user entered the correct password and was accepted by the SSH daemon, but the session was closed immediately due to the invalid shell. The other options are incorrect because they are not supported by the outputs. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 413.

NEW QUESTION: 46
A Linux administrator is creating a primary partition on the replacement hard drive for an application server.
Which of the following commands should the administrator issue to verify the device name of this partition?
A. sudo fdisk /dev/sda
B. sudo fdisk -s /dev/sda
C. sudo fdisk -l
D. sudo fdisk -h
Answer: C (LEAVE A REPLY)
Explanation
The command sudo fdisk -l should be issued to verify the device name of the partition. The sudo command allows the administrator to run commands as the superuser or another user. The fdisk command is a tool for manipulating disk partitions on Linux systems. The -l option lists the partitions on all disks or a specific disk.
The command sudo fdisk -l will show the device names, sizes, types, and other information of the partitions on all disks. The administrator can identify the device name of the partition by looking at the output. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not list the partitions (sudo fdisk /dev/sda or sudo fdisk -h) or do not exist (sudo fdisk -s /dev/sda). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 317.

NEW QUESTION: 47
A Linux administrator is troubleshooting SSH connection issues from one of the workstations.
When users attempt to log in from the workstation to a server with the IP address 104.21.75.76, they receive the following message:

```
ssh: connect to host 104.21.75.76 port 22: Connection refused
```

The administrator reviews the information below:

Workstation output 1:

```
eth0: <BROADCAST,MULTICAST, UP, LOWER_UP> mtu 1500 qdisc mq state UP group default
link/ether 00:15:5d:e9:e9:fb brd 5.189.153.255 scope global eth0
inet 5.189.153.89/24 brd 5.189.153.255 scope global eth0
```

Workstation output 2:

```
default via 5.189.153.1 dev eth0
5.189.153.0/24 dev eth0 proto kertnel scope link src 5.189.153.89
```

Server output 1:

```
target     prot   opt   source            destination
REJECT     tcp    --    101.68.78.194     0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                       reject-with icmp-port-unreachable
REJECT     tcp    --    222.186.180.130   0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                       reject-with icmp-port-unreachable
REJECT     tcp    --    104.131.1.39      0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                       reject-with icmp-port-unreachable
REJECT     tcp    --    68.183.196.11     0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                       reject-with icmp-port-unreachable
REJECT     tcp    --    5.189.153.89      0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                       reject-with icmp-port-unreachable
REJECT     tcp    --    41.93.32.148      0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                       reject-with icmp-port-unreachable
```

Server output 2:

```
sshd. service - OpenSSH server daemon
  Loaded: loaded (/usr/lib/systemd/system/sshd.service: disabled: vendor preset: enabled)
  Active: active (running) since Thu 2021-08-26 18:50:19 CEST; 2 weeks 5 days ago
```

Server output 3:

```
eth0: <BROADCAST, MULTICAST, UP, LOWER_UP> mtu 1500 qdisc mg state UP group default
link/ether 52:52:00:2a:bb:98 brd 104.21.75.255 scope global eth0
inet 104.21.75.76/24 brd 104.21.75.255 scope qlobal eth0
```

Server output 4:

```
default via 104.21.75.254 dev eth0
104.21.75.0/24 dev eth0 proto kertnel scope link src 104.21.75.76
```

Which of the following is causing the connectivity issue?

A. The workstation has the wrong IP settings.

B. The sshd service is disabled.

C. The server's firewall is preventing connections from being made.
D. The server has an incorrect default gateway configuration.
Answer: C (LEAVE A REPLY)
The server's firewall is preventing connections from being made, which is causing the connectivity issue. The output of iptables -L -n shows that the firewall is blocking all incoming traffic on port 22, which is the default port for SSH. The output of ssh -v user@104.21.75.76 shows that the connection is refused by the server. To resolve the issue, the administrator needs to allow port 22 on the firewall. The other options are incorrect because they are not supported by the outputs. The workstation has the correct IP settings, as shown by the output of ip addr show. The sshd service is enabled and running, as shown by the output of systemctl status sshd. The server has the correct default gateway configuration, as shown by the output of ip route show. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, pages 406-407.

NEW QUESTION: 48
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 (LEAVE A REPLY)
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: 49
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 $(docker ps -aq)
B. docker rm --all
C. docker rm --state exited
D. docker images prune *
Answer: C (LEAVE A REPLY)

NEW QUESTION: 50

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.sh
chmod 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 (LEAVE A REPLY)

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). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing File Permissions and Ownership, pages 346-348.

NEW QUESTION: 51

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

A. The checkdiskspace.service is not running.

B. The checkdiskspace.service needs to be enabled.

C. The OnCalendar schedule is incorrect in the timer definition.

D. The system-daemon services need to be reloaded.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 52
Due to low disk space, a Linux administrator finding and removing all log files that were modified more than 180 days ago. Which of the following commands will accomplish this task?
A. find /var/log -type d -mtime +180 -print -exec rm {} \;
B. find /var/log -type f -modified +180 -rm
C. find /var/log -type f -mtime +180 -exec rm {} \
D. find /var/log -type c -atime +180 -remove
Answer: (SHOW ANSWER)
The command that will accomplish the task of finding and removing all log files that were modified more than 180 days ago is find /var/log -type f -mtime +180 -exec rm {} ;. This command will use find to search for files (-type f) under /var/log directory that have a modification time (-mtime) older than 180 days (+180). For each matching file, it will execute (-exec) the rm command to delete it, passing the file name as an argument ({}). The command will end with a semicolon (;), which is escaped with a backslash to prevent shell interpretation.
The other options are not correct commands for accomplishing the task. The find /var/log -type d -mtime +180 -print -exec rm {} ; command will search for directories (-type d) instead of files, and print their names (-print) before deleting them. This is not what the task requires.
The find /var/log -type f -modified +180 -rm command is invalid because there is no such option as -modified or -rm for find. The correct options are -mtime and -delete, respectively. The find /var/log -type c -atime +180 -remove command is also invalid because there is no such option as -remove for find. Moreover, it will search for character special files (-type c) instead of regular files, and use access time (-atime) instead of modification time. Reference: find(1) - Linux manual page; Find and delete files older than n days in Linux

NEW QUESTION: 53
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      9.86      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    25.48    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 (LEAVE A REPLY)
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: 54
A user is asking the systems administrator for assistance with writing a script to verify whether a file exists. Given the following:

```
#1/bin/bash
filename=$1
<CONDITIONAL>
echo "File exists"
else
echo "File does not exist"
fi
```


Which of the following commands should replace the <CONDITIONAL> string?
A. if [-f "$filename"]; then
B. if [-d "$filename"]; then
C. if [-f "$filename"] then
D. if [-f "$filename"]; while
Answer: A (LEAVE A REPLY)
The command if [-f "$filename"]; then checks if the variable $filename refers to a regular file that exists. The -f option is used to test for files. If the condition is true, the commands after then are executed. This is the correct way to replace the <CONDITIONAL> string. The other options are incorrect because they either use the wrong option (-d tests for directories), the wrong syntax (missing a semicolon after

the condition), or the wrong keyword (while is used for loops, not conditions). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Writing and Executing Bash Shell Scripts, page 493.

NEW QUESTION: 55
A junior administrator is trying to set up a passwordless SSH connection to one of the servers. The administrator follows the instructions and puts the key in the authorized_key file at the server, but the administrator is still asked to provide a password during the connection. Given the following output:

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

Which of the following commands would resolve the issue and allow an SSH connection to be established without a password?
A. restorecon -rv .ssh/authorized_key
B. mv .ssh/authorized_key .ssh/authorized_keys
C. systemct1 restart sshd.service
D. chmod 600 mv .ssh/authorized_key
Answer: (SHOW ANSWER)
Explanation
The command mv .ssh/authorized_key .ssh/authorized_keys will resolve the issue and allow an SSH connection to be established without a password. The issue is caused by the incorrect file name of the authorized key file on the server. The file should be named authorized_keys, not authorized_key.
The mv command will rename the file and fix the issue. The other options are incorrect because they either do not affect the file name (restorecon or chmod) or do not restart the SSH service (systemct1). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 410.

NEW QUESTION: 56
In order to copy data from another VLAN, a systems administrator wants to temporarily assign IP address 10.0.6 5/24 to the newly added network interface enp1s0f1. Which of the following commands should the administrator run to achieve the goal?
A. ifconfig 10.0.6.5/24 enpsIs0f1
B. echo "IPV4_ADDRESS=10.0.6.5/24" > /etc/sysconfig/network-scripts/ifcfg-enplsOfl
C. ip addr add 10.0.6.5/24 dev enpls0f1
D. nmcli conn add lpv4.address-10.0.6.5/24 ifname enpls0f1
Answer: C (LEAVE A REPLY)

NEW QUESTION: 57
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 (LEAVE A REPLY)

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

NEW QUESTION: 58
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 (LEAVE A REPLY)
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: 59
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 (LEAVE A REPLY)
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: 60
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: (SHOW ANSWER)
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: 61
Several users reported that they were unable to write data to the /oracle1 directory. The following output has been provided:

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

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

A. df -i /oracle1

B. fdisk -1 /dev/sdb1

C. lsblk /dev/sdb1

D. du -sh /oracle1

Answer: (SHOW ANSWER)

Explanation

The administrator should use the command df -i /oracle1 to diagnose the issue of users being unable to write data to the /oracle1 directory. This command will show the inode usage of the /oracle1 filesystem, which indicates how many files and directories can be created on it. If the inode usage is 100%, it means that no more files or directories can be added, even if there is still free space on the disk. The administrator can then delete some unnecessary files or directories, or increase the inode limit of the filesystem, to resolve the issue.

The other options are not correct commands for diagnosing this issue. The fdisk -l /dev/sdb1 command will show the partition table of /dev/sdb1, which is not relevant to the inode usage. The lsblk /dev/sdb1 command will show information about /dev/sdb1 as a block device, such as its size, mount point, and type, but not its inode usage. The du -sh /oracle1 command will show the disk usage of /oracle1 in human-readable format, but not its inode usage. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 7: Managing Disk Storage; How to Check Inode Usage in Linux - Fedingo

NEW QUESTION: 62

A Linux administrator is troubleshooting SSH connection issues from one of the workstations.

When users attempt to log in from the workstation to a server with the IP address 104.21.75.76, they receive the following message:

```
ssh: connect to host 104.21.75.76 port 22: Connection refused
```

The administrator reviews the information below:

Workstation output 1:

```
eth0: <BROADCAST,MULTICAST, UP, LOWER_UP> mtu 1500 qdisc mq state UP group default
link/ether 00:15:5d:e9:e9:fb brd 5.189.153.255 scope global eth0
inet 5.189.153.89/24 brd 5.189.153.255 scope global eth0
```

Workstation output 2:

```
default via 5.189.153.1 dev eth0
5.189.153.0/24 dev eth0 proto kertnel scope link src 5.189.153.89
```

Server output 1:

```
target     prot   opt   source              destination
REJECT     tcp    --    101.68.78.194       0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                         reject-with icmp-port-unreachable
REJECT     tcp    --    222.186.180.130     0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                         reject-with icmp-port-unreachable
REJECT     tcp    --    104.131.1.39        0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                         reject-with icmp-port-unreachable
REJECT     tcp    --    68.183.196.11       0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                         reject-with icmp-port-unreachable
REJECT     tcp    --    5.189.153.89        0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                         reject-with icmp-port-unreachable
REJECT     tcp    --    41.93.32.148        0.0.0.0/0    tcp dpt:22 ctstate NEW, UNTRACKED
                                                         reject-with icmp-port-unreachable
```

Server output 2:

```
sshd. service - OpenSSH server daemon
  Loaded: loaded (/usr/lib/systemd/system/sshd.service: disabled: vendor preset: enabled)
  Active: active (running) since Thu 2021-08-26 18:50:19 CEST; 2 weeks 5 days ago
```

Server output 3:

```
eth0: <BROADCAST, MULTICAST, UP, LOWER_UP> mtu 1500 qdisc mg state UP group default
link/ether 52:52:00:2a:bb:98 brd 104.21.75.255 scope global eth0
inet 104.21.75.76/24 brd 104.21.75.255 scope global eth0
```

Server output 4:

```
default via 104.21.75.254 dev eth0
104.21.75.0/24 dev eth0 proto kertnel scope link src 104.21.75.76
```

Which of the following is causing the connectivity issue?

A. The workstation has the wrong IP settings.

B. The sshd service is disabled.

C. The server's firewall is preventing connections from being made.

D. The server has an incorrect default gateway configuration.

Answer: (SHOW ANSWER)

Explanation

The server's firewall is preventing connections from being made, which is causing the connectivity issue. The output of iptables -L -n shows that the firewall is blocking all incoming traffic on port 22, which is the default port for SSH. The output of ssh -v user@104.21.75.76 shows that the connection is refused by the server. To resolve the issue, the administrator needs to allow port 22 on the firewall. The other options are incorrect because they are not supported by the outputs. The workstation has the correct IP settings, as shown by the output of ip addr show. The sshd service is enabled and running, as shown by the output of systemctl status sshd. The server has the correct default gateway configuration, as shown by the output of ip route show. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, pages 406-407.

NEW QUESTION: 63

A new Linux systems administrator just generated a pair of SSH keys that should allow connection to the servers. Which of the following commands can be used to copy a key file to remote servers? (Choose two.)

A. wget

B. ssh-keygen

C. ssh-keyscan

D. ssh-copy-id

E. ftpd

F. scp

Answer: D,F (LEAVE A REPLY)

Explanation

The commands ssh-copy-id and scp can be used to copy a key file to remote servers. The command ssh-copy-id copies the public key to the authorized_keys file on the remote server, which allows the user to log in without a password. The command scp copies files securely over SSH, which can be used to transfer the key file to any location on the remote server. The other options are incorrect because they are not related to copying key files. The command wget downloads files from the web, the command ssh-keygen generates key pairs, the command ssh-keyscan collects public keys from remote hosts, and the command ftpd is a FTP server daemon. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, pages 408-410.

NEW QUESTION: 64

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

Routing table:

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

IP configuration:

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

ARP table:

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

```
/etc/resolv.conf:
Generated by NetworkHanager
search company.com
nameserver 10.0.5.1
```


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

A. An internal-only DNS server is configured.

B. The IP netmask is wrong for ens3.

C. Two default routes are configured.

D. The ARP table contains incorrect entries.

Answer: C (LEAVE A REPLY)

Explanation

The most likely cause of the issue is that two default routes are configured on the server. The default route is the route that is used when no other route matches the destination of a packet. The default route is usually the gateway that connects the local network to the Internet. The server configuration shows that there are two default routes in the routing table, one with the gateway 192.168.1.1 and the other with the gateway 10.0.0.1.

This can cause a conflict and confusion for the server when deciding which gateway to use for the outgoing packets. Some packets may be sent to the wrong gateway and fail to reach the Internet, while some packets may be sent to the correct gateway and work as intended. This can result in intermittent connection issues and inconsistent behavior. The administrator should remove one of the default routes and keep only the correct one for the network. This can be done by using the ip route del command or by editing the network configuration files. This will resolve the issue and restore the connectivity. The other options are incorrect because they are not supported by the

outputs. The DNS server, the IP netmask, and the ARP table are not the causes of the issue. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, pages 381-382.

NEW QUESTION: 65
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: (SHOW ANSWER)
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: 66
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	6[illegible]	R	0.3	1.1	00:17.6	uploadpw.sh
4846	joe	30	-10	6426	6[illegible]	7[illegible]	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 (LEAVE A REPLY)
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: 67

A systems administrator is investigating an issue in which one of the servers is not booting up properly. The journalctl entries show the following:

```
Sep 16 20:30:43 server kernel: acpi PNP0A03:00: _OSC failed (AE_NOT_FOUND);
-- Subject: Unit dev-mapper-centos\x2dapp.device has failed
Sep 16 20:32:15 server systemd[1]: Dependency failed for /opt/app
-- Subject: Unit opt-app.mount has failed
-- Unit opt-app.mount has failed
Sep 16 20:32:15 server systemd[1]: Dependency failed for Local File Systems.
-- Subject: Unit local-fs.target has failed
-- Unit local-fs.target has failed.
Sep 16 20:32:15 server systemd[1]: Dependency failed for Relabel all filesystem, if necessary.
-- Subject: Unit rhel-autorelabel.service has failed
-- Unit rhel-autorelabel.service has failed.
```

Which of the following will allow the administrator to boot the Linux system to normal mode quickly?

A. Comment out the /opt/app filesystem in /etc/fstab and reboot.

B. Reformat the /opt/app filesystem and reboot.

C. Perform filesystem checks on local filesystems and reboot.

D. Trigger a filesystem relabel and reboot.

Answer: A (LEAVE A REPLY)

Explanation

The fastest way to boot the Linux system to normal mode is to comment out the /opt/app filesystem in /etc/fstab and reboot. This will prevent the system from trying to mount the /opt/app filesystem at boot time, which causes an error because the filesystem does not exist or is corrupted. Commenting out a line in /etc/fstab can be done by adding a # symbol at the beginning of the line. Rebooting the system will apply the changes and allow the system to boot normally. Reformatting the /opt/app filesystem will not help to boot the system, as it will erase any data on the filesystem and require manual intervention to create a new filesystem.

Performing filesystem checks on local filesystems will not help to boot the system, as it will not fix the missing or corrupted /opt/app filesystem. Triggering a filesystem relabel will not help to boot the system, as it will only change the security context of files and directories according to SELinux policy. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Disk Storage, page 456.

NEW QUESTION: 68

A systems engineer is adding a new 1GB XFS filesystem that should be temporarily mounted under /ops/app.

Which of the following is the correct list of commands to achieve this goal?

A.

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

B.

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

C.

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

D.

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

Answer: D (LEAVE A REPLY)

Explanation

The list of commands in option D is the correct way to achieve the goal. The commands are as follows:

fallocate -l 1G /ops/app.img creates a 1GB file named app.img under the /ops directory.

mkfs.xfs /ops/app.img formats the file as an XFS filesystem.

mount -o loop /ops/app.img /ops/app mounts the file as a loop device under the /ops/app directory. The other options are incorrect because they either use the wrong commands (dd or truncate instead of fallocate), the wrong options (-t or -f instead of -o), or the wrong order of arguments (/ops/app.img

/ops/app instead of /ops/app /ops/app.img). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 323-324.

NEW QUESTION: 69

A Linux administrator booted up the server and was presented with a non-GUI terminal. The administrator ran the command systemctl isolate graphical.target and rebooted the system by running systemctl reboot, which fixed the issue. However, the next day the administrator was presented again with a non-GUI terminal. Which of the following is the issue?

A. The administrator did not reboot the server properly.

B. The administrator did not set the default target to basic.target.

C. The administrator did not set the default target to graphical.target.

D. The administrator did not shut down the server properly.

Answer: (SHOW ANSWER)

The issue is that the administrator did not set the default target to graphical.target. A target is a unit of systemd that groups together other units by a common purpose or state. The graphical.target is a target that starts the graphical user interface (GUI) along with other services. The administrator used the command systemctl isolate graphical.target to switch to this target temporarily, but this does not change the default target that is activated at boot time. To make this change permanent, the administrator should have used the command systemctl

set-default graphical.target, which creates a symbolic link from /etc/systemd/system/default.target to /usr/lib/systemd/system/graphical.target.
The other options are not correct explanations for the issue. The administrator did reboot the server properly by using systemctl reboot, which shuts down and restarts the system cleanly. The administrator did not need to set the default target to basic.target, which is a minimal target that only starts essential services. The administrator did not shut down the server improperly, which could have caused file system corruption or data loss, but not affect the default target. Reference: systemctl(1) - Linux manual page; How to Change Runlevels (targets) in SystemD

NEW QUESTION: 70
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 POSTROUTING -o eth0 -j MASQUERADE
```

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

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

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

Answer: B (LEAVE A REPLY)

NEW QUESTION: 71
A Linux administrator needs to resolve a service that has failed to start. The administrator runs the following command:

```
ls -l startup file
```

The following output is returned

```
----------. root root 81k Sep 13 19:01 startupfile
```

Which of the following is MOST likely the issue?
A. The service does not have permissions to read write the startupfile.
B. The service startupfile size cannot be 81k.
C. The service startupfile cannot be owned by root.

D. The service startupfile should not be owned by the root group.
Answer: (SHOW ANSWER)
Explanation
The most likely issue is that the service does not have permissions to read or write the startupfile. The output of systemct1 status startup.service shows that the service has failed to start and the error message is
"Permission denied". The output of ls -l /etc/startupfile shows that the file has the permissions -rw-r--r--, which means that only the owner (root) can read and write the file, while the group (root) and others can only read the file. The service may not run as root and may need write access to the file. The administrator should change the permissions of the file by using the chmod command and grant write access to the group or others, or change the owner or group of the file by using the chown command and assign it to the user or group that runs the service. The other options are incorrect because they are not supported by the outputs. The file size, owner, and group are not the causes of the issue. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, pages 345-346.

NEW QUESTION: 72
Users have been unable to reach www.comptia.org from a Linux server. A systems administrator is troubleshooting the issue and does the following:

Output 1:

```
2: eth0: <BROADCAST,MULTICAST,UP, LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
   link/ether ac:11:22:33:44:cd brd ff:ff:ff:ff:ff:ff
   inet 192.168.168.10/24 brd 192.168.169.255 scope global dynamic noprefixroute eth0
      valid_lft 8097sec preferred_lft 8097sec
   inet fe80::4daf:8c7c:a6ff:2771/64 scope link noprefixroute
      valid_lft forever preferred_lft forever
```

Output 2:

```
nameserver 192.168.168.53
```

Output 3:

```
FING 192.168.168.53 (192.168.168.53) 56(84) bytes of data.
64 bytes from 192.168.168.53: icmp_seq=1 ttl=64 time=2.85 ms

--- 192.168.168.53 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 2.847/2.847/2.847/0.000 ms
```

Output 4:

```
192.168.168.0/24 dev eth0 proto kernel scope link src 192.168.168.10 metric 600
```

Output 5:

```
...
;; QUESTION SECTION:
;www.comptia.org. IN A

;; ANSWER SECTION:
. 0 CLASS4096 OPT 10 8 LgmNvk0AazU=

;; ADDITIONAL SECTION:
www.comptia.org. 3385 IN A 23.96.239.26
...
```

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

A. The name www.comptia.org does not point to a valid IP address.

B. The server 192.168.168.53 is unreachable.

C. No default route is set on the server.

D. The network interface eth0 is disconnected.

Answer: B (LEAVE A REPLY)

Explanation

The issue is caused by the server 192.168.168.53 being unreachable. This server is the DNS server configured in the /etc/resolv.conf file, which is used to resolve domain names to IP addresses. The ping command shows that the server cannot be reached, and the nslookup command shows that the name www.comptia.org cannot be resolved using this server. The other options are incorrect because:

The name www.comptia.org does point to a valid IP address, as shown by the nslookup command using another DNS server (8.8.8.8).

The default route is set on the server, as shown by the ip route command, which shows a default gateway of 192.168.168.1.

The network interface eth0 is connected, as shown by the ip link command, which shows a state of UP for eth0. References: CompTIA Linux+ Study Guide, Fourth Edition, page 457-458, 461-462.

NEW QUESTION: 73
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 application is listening on the loopback interface.
B. The IP address 0.0.0.0 is not valid.
C. The application is listening on port 1234.
D. The application is not running.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 74
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. rsync
C. netstat
D. host
Answer: A (LEAVE A REPLY)
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). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 378.

NEW QUESTION: 75
An administrator attempts to rename a file on a server but receives the following error.

```
mv: cannot move 'files/readme.txt' to 'files/readme.txt.orig': Operation not permitted.
```

The administrator then runs a few commands and obtains the following output:

```
$ ls -ld files/
drwxrwxrwt.1    users    users    20    Sep 10 15:15    files/
$ ls -a files/
drwxrwxrwt.1    users    users    20    Sep 10 15:15    .
drwxr-xr-x.1    users    users    32    Sep 10 15:15    ..
-rw-rw-r--.1    users    users    4     Sep 12 10:34    readme.txt
```

Which of the following commands should the administrator run NEXT to allow the file to be renamed by any user?

A. chgrp reet files

B. chacl -R 644 files

C. chown users files

D. chmod -t files

Answer: (SHOW ANSWER)

Explanation

The command that the administrator should run NEXT to allow the file to be renamed by any user is chmod -t files. This command uses the chmod tool, which is used to change file permissions and access modes. The -t option removes (or sets) the sticky bit on a directory, which restricts deletion or renaming of files within that directory to only their owners or root. In this case, since files is a directory with sticky bit set (indicated by t in drwxrwxrwt), removing it will allow any user to rename or delete files within that directory.

The other options are not correct commands for allowing any user to rename files within files directory. The chgrp reet files command will change the group ownership of files directory to reet, but it will not affect its permissions or access modes. The chacl -R 644 files command is invalid, as chacl is used to change file access control lists (ACLs), not permissions or access modes. The chown users files command will change the user ownership of files directory to users, but it will not affect its permissions or access modes. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing Users and Groups; chmod(1) - Linux manual page

NEW QUESTION: 76

After installing some RPM packages, a systems administrator discovers the last package that was installed was not needed. Which of the following commands can be used to remove the package?

A. dnf remove packagename

B. apt-get remove packagename
C. rpm -i packagename
D. apt remove packagename
Answer: A (LEAVE A REPLY)
Explanation
The command that can be used to remove an RPM package that was installed by mistake is dnf remove packagename. This command will use the DNF package manager to uninstall an RPM package and its dependencies from a Linux system that uses RPM-based distributions, such as Red Hat Enterprise Linux or CentOS. The DNF package manager handles dependency resolution and metadata searching for RPM packages.
The other options are not correct commands for removing an RPM package from a Linux system. The apt-get remove packagename and apt remove packagename commands are used to remove Debian packages from a Linux system that uses Debian-based distributions, such as Ubuntu or Debian. They are not compatible with RPM packages. The rpm -i packagename command is used to install an RPM package, not to remove it. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 9: Managing Software Packages; How to install/remove/query/update RPM packages in Linux (Cheat Sheet ...

NEW QUESTION: 77
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
```

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

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

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

D.

Answer: C (LEAVE A REPLY)

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: 78

A systems administrator is investigating an issue in which one of the servers is not booting up properly. The journalctl entries show the following:

```
Sep 16 20:30:43 server kernel: acpi PNP0A03:00: _OSC failed (AE_NOT_FOUND);
-- Subject: Unit dev-mapper-centos\x2dapp.device has failed
Sep 16 20:32:15 server systemd[1]: Dependency failed for /opt/app
-- Subject: Unit opt-app.mount has failed
-- Unit opt-app.mount has failed
Sep 16 20:32:15 server systemd[1]: Dependency failed for Local File Systems.
-- Subject: Unit local-fs.target has failed
-- Unit local-fs.target has failed.
Sep 16 20:32:15 server systemd[1]: Dependency failed for Relabel all filesystem, if necessary.
-- Subject: Unit rhel-autorelabel.service has failed
-- Unit rhel-autorelabel.service has failed.
```

Which of the following will allow the administrator to boot the Linux system to normal mode quickly?

A. Comment out the /opt/app filesystem in /etc/fstab and reboot.

B. Reformat the /opt/app filesystem and reboot.

C. Perform filesystem checks on local filesystems and reboot.

D. Trigger a filesystem relabel and reboot.

Answer: A (LEAVE A REPLY)

Explanation

The fastest way to boot the Linux system to normal mode is to comment out the /opt/app filesystem in /etc/fstab and reboot. This will prevent the system from trying to mount the /opt/app filesystem at boot time, which causes an error because the filesystem does not exist or is corrupted. Commenting out a line in /etc/fstab can be done by adding a # symbol at the beginning of the line. Rebooting the system will apply the changes and allow the system to boot normally. Reformatting the /opt/app filesystem will not help to boot the system, as it will erase any data on the filesystem and require manual intervention to create a new filesystem.

Performing filesystem checks on local filesystems will not help to boot the system, as it will not fix the missing or corrupted /opt/app filesystem. Triggering a filesystem relabel will not help to boot the system, as it will only change the security context of files and directories according to SELinux policy. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Disk Storage, page 456.

NEW QUESTION: 79
A systems administrator is troubleshooting connectivity issues and trying to find out why a Linux server is not able to reach other servers on the same subnet it is connected to. When listing link parameters, the following is presented:

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


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

NEW QUESTION: 80
Joe, a user, is unable to log in to the Linux system Given the following output:

```
# grep joe /etc/passwd /etc/shadow
/etc/passwd:joe:x:1001:1001::/home/joe:/bin/nologin
/etc/shadow:joe:$6$3uOw6qWx9876jGhgKJsdfH987634534voj.:18883:0:99999:7:::
```

Which of the following command would resolve the issue?
A. passwd -u joe
B. usermod -s /bin/bash joe
C. chage -E 90 joe
D. pam_tally2 -u joe -r
Answer: B (LEAVE A REPLY)

NEW QUESTION: 81
A systems administrator pressed Ctrl+Z after starting a program using the command line, and the shell prompt was presented. In order to go back to the program, which of the following commands can the administrator use?
A. fg
B. su

C. bg
D. ed
Answer: (SHOW ANSWER)
Ctrl+Z suspended the process, and "fg" will bring it back into the foreground of the shell

NEW QUESTION: 82
A systems administrator detected corruption in the /data filesystem. Given the following output:

```
root@localhost ~]# lsblk -f
```

NAME	FSTYPE	LABEL/UUID	MOUNTPOINT
sda			
├─sda1	vfat	4E7D-9539	/boot/efi
├─sda2	xfs	98442caf-473d-448e-aee5-561a82297314	/boot
├─sda3	swap	19f064e4-7c51-4b02-8219-99362a3c45ec	[SWAP]
├─sda4	xfs	25d96ada-4289-4def-9202-6ab11affbed3	/
├─sda5	xfs	61435ee9-855d-4de9-9c67-39aeb7f3edb5	/home
sdc			
├─sdc1	ext4	92435ff9-745e-4fg9-9c67-39aeb7f3exf5	/data

Which of the following commands can the administrator use to best address this issue?
A. umount /data
mkfs . xfs /dev/sclcl
mount /data
B. umount /data
xfs repair /dev/ sdcl
mount /data

C. umount /data
fsck /dev/ sdcl
mount / data
D. umount /data
pvs /dev/sdcl
mount /data
Answer: B (LEAVE A REPLY)
Explanation
The xfs repair command is used to check and repair an XFS filesystem, which is the type of filesystem used for the /data partition, as shown in the output. The administrator needs to unmount the /data partition before running the xfs repair command on it, and then mount it back after the repair is done. For example: umount
/data; xfs_repair /dev/sdcl; mount /data. The mkfs.xfs command is used to create a new XFS filesystem, which would erase all the data on the partition. The fsck command is used to check and repair other types of filesystems, such as ext4, but not XFS. The pvs command is used to display information about physical volumes in a logical volume manager (LVM) setup, which is not relevant for this issue.

NEW QUESTION: 83
A Linux administrator needs to resolve a service that has failed to start. The administrator runs the following command:

```
ls -l startup file
```

The following output is returned

```
----------. root root 81k Sep 13 19:01 startupfile
```

Which of the following is MOST likely the issue?
A. The service does not have permissions to read write the startupfile.
B. The service startupfile size cannot be 81k.
C. The service startupfile cannot be owned by root.
D. The service startupfile should not be owned by the root group.
Answer: A (LEAVE A REPLY)
The most likely issue is that the service does not have permissions to read or write the startupfile. The output of systemctl status startup.service shows that the service has failed to start and the error message is "Permission denied". The output of ls -l /etc/startupfile shows that the file has the permissions -rw-r--r--, which means that only the owner (root) can read and write the file, while the group (root) and others can only read the file. The service may not run as root and may need write access to the file. The administrator should change the permissions of the file by using the chmod command and grant write access to the group or others, or change the owner or group of the file by using the chown command and assign it to the user or group that runs the service. The other options are incorrect because they are not supported by the outputs. The file size, owner, and group are not the causes of the issue. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, pages 345-346.

NEW QUESTION: 84
A Linux user reported the following error after trying to connect to the system remotely:

ssh: connect to host 10.0.1.10 port 22: Resource temporarily unavailable The Linux systems administrator executed the following commands in the Linux system while trying to diagnose this issue:

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

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


Which of the following commands will resolve this issue?

A. firewall-cmd --zone=public --permanent --add-service=22

B. systemct1 enable firewalld; systemct1 restart firewalld

C. firewall-cmd --zone=public --permanent --add-service=ssh

D. firewall-cmd --zone=public --permanent --add-port=22/udp

Answer: C (LEAVE A REPLY)

Explanation

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

NEW QUESTION: 85

A developer reported an incident involving the application configuration file /etc/httpd/conf/httpd.conf that is missing from the server. Which of the following identifies the RPM package that installed the configuration file?

A. rpm -qf /etc/httpd/conf/httpd.conf
B. rpm -ql /etc/httpd/conf/httpd.conf
C. rpm -query /etc/httpd/conf/httpd.conf
D. rpm -q /etc/httpd/conf/httpd.conf
Answer: A (LEAVE A REPLY)
Explanation
The rpm -qf /etc/httpd/conf/httpd.conf command will identify the RPM package that installed the configuration file. This command will query the database of installed packages and display the name of the package that owns the specified file. The rpm -ql /etc/httpd/conf/httpd.conf command is invalid, as -ql is not a valid option for rpm. The rpm --query /etc/httpd/conf/httpd.conf command is incorrect, as --query requires a package name, not a file name. The rpm -q /etc/httpd/conf/httpd.conf command is incorrect, as -q requires a package name, not a file name. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter
19: Managing Packages and Software, page 560.

NEW QUESTION: 86
A Linux engineer has been notified about the possible deletion of logs from the file /opt/app/logs. The engineer needs to ensure the log file can only be written into without removing previous entries.

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

Which of the following commands would be BEST to use to accomplish this task?
A. chattr +d /opt/app/logs
B. chattr +i /opt/app/logs
C. chattr +a /opt/app/logs
D. chattr +c /opt/app/logs
Answer: C (LEAVE A REPLY)

NEW QUESTION: 87
A Linux administrator needs to ensure that Java 7 and Java 8 are both locally available for developers to use when deploying containers. Currently only Java 8 is available. Which of the following commands should the administrator run to ensure both versions are available?
A. docker image load java:7
B. docker image pull java:7
C. docker image import java:7
D. docker image build java:7
Answer: B (LEAVE A REPLY)
Explanation
The command that the administrator should run to ensure that both Java 7 and Java 8 are locally available for developers to use when deploying containers is docker image pull java:7. This command will use the docker image pull subcommand to download the java:7 image from Docker Hub, which is the default registry for Docker images. The java:7 image contains Java 7 installed on a Debian-based Linux system. The administrator can also specify a different registry by using the syntax registry/repository:tag.

The other options are not correct commands for ensuring that both Java 7 and Java 8 are locally available for developers to use when deploying containers. The docker image load java:7 command will load an image from a tar archive or STDIN, not from a registry. The docker image import java:7 command will create a new filesystem image from the contents of a tarball, not from a registry. The docker image build java:7 command will build an image from a Dockerfile, not from a registry. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks; docker image pull | Docker Docs

NEW QUESTION: 88
Which of the following technologies can be used as a central repository of Linux users and groups?
A. LDAP
B. MFA
C. SSO
D. PAM
Answer: (SHOW ANSWER)
Explanation
LDAP stands for Lightweight Directory Access Protocol, which is a protocol for accessing and managing a central directory of users and groups. LDAP can be used as a central repository of Linux users and groups, allowing for centralized authentication and authorization across multiple Linux systems. MFA, SSO, and PAM are not technologies that can be used as a central repository of Linux users and groups. MFA stands for Multi-Factor Authentication, which is a method of verifying a user's identity using more than one factor, such as a password, a token, or a biometric. SSO stands for Single Sign-On, which is a feature that allows a user to log in once and access multiple applications or systems without having to re-enter credentials. PAM stands for Pluggable Authentication Modules, which is a framework that allows Linux to use different authentication methods, such as passwords, tokens, or biometrics. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing Users and Groups

NEW QUESTION: 89
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 /.

Answer:

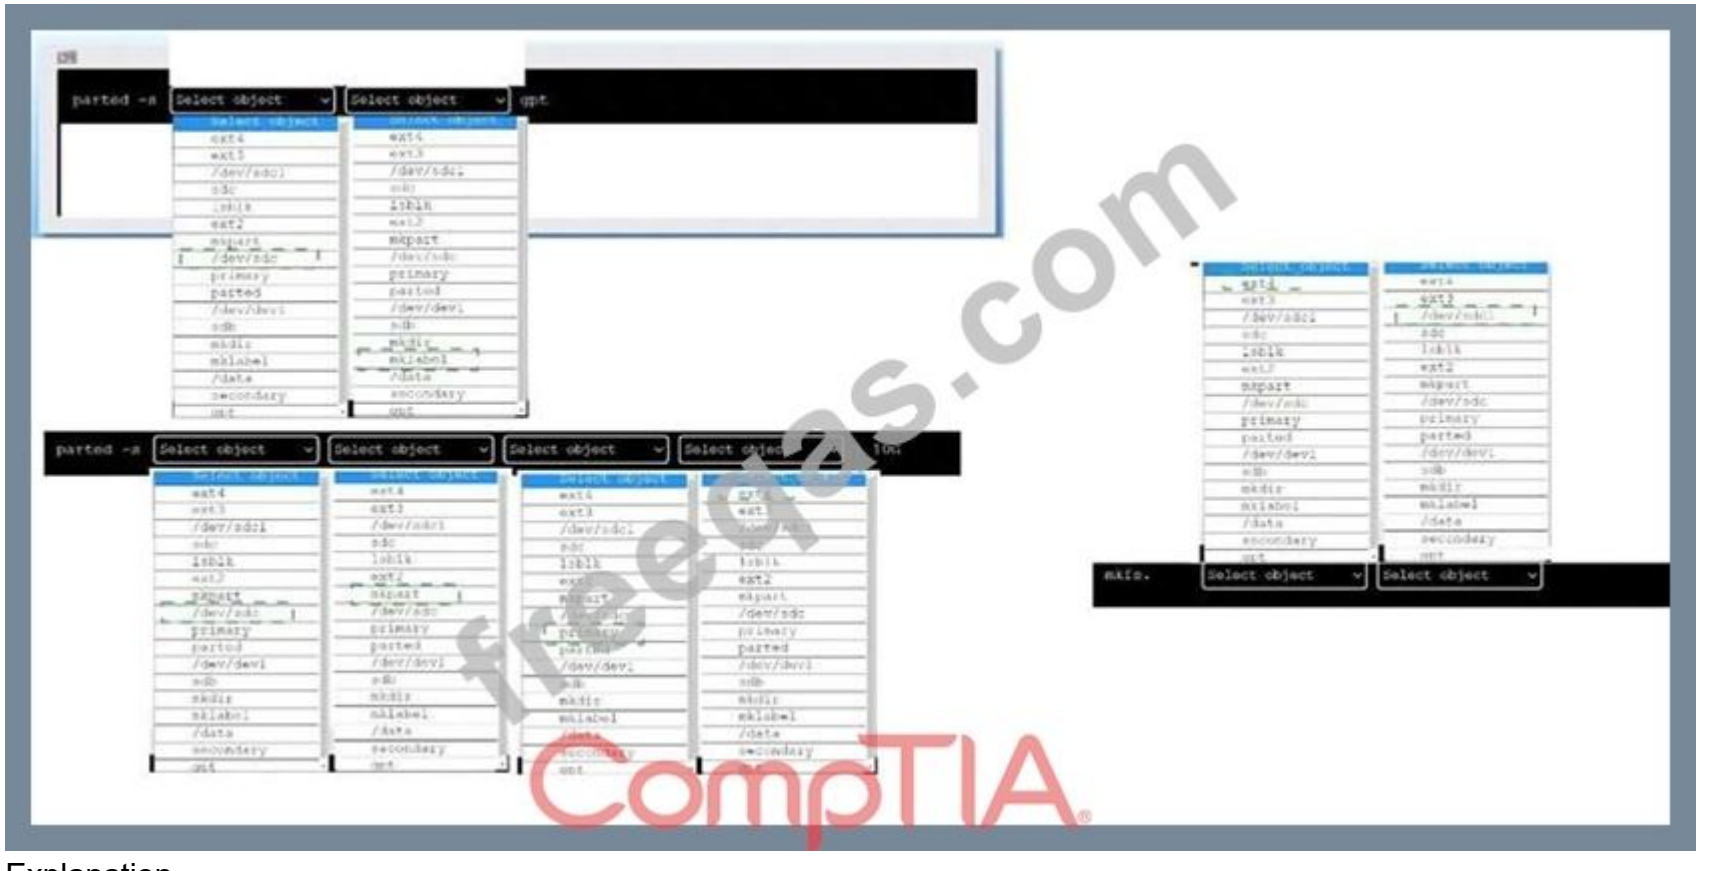

Explanation

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: 90

A junior Linux administrator is tasked with installing an application. The installation guide states the application should only be installed in a run level 5 environment.

```
$ systemctl get-default
getty.target
```

Which of the following commands would ensure the server is set to runlevel 5?

A. systemct1 isolate multi-user.target

B. systemct1 isolate graphical.target
C. systemct1 isolate network.target
D. systemct1 isolate basic.target
Answer: (SHOW ANSWER)
Explanation
The command that would ensure the server is set to runlevel 5 is systemct1 isolate graphical.target. This command will change the current target (or runlevel) of systemd to graphical.target, which is equivalent to runlevel 5 in SysV init systems. Graphical.target means that the system will start with a graphical user interface (GUI) and all services required for it.
The other options are not correct commands for setting the server to runlevel 5. The systemct1 isolate multi-user.target command will change the current target to multi-user.target, which is equivalent to runlevel 3 in SysV init systems. Multi-user.target means that the system will start with multiple user logins and networking, but without a GUI. The systemct1 isolate network.target command will change the current target to network.target, which is not a real runlevel but a synchronization point for network-related services.
Network.target means that network functionality should be available, but does not specify whether it should be started before or after it.
The systemct1 isolate basic.target command will change the current target to basic.target, which is also not a real runlevel but a synchronization point for basic system services. Basic.target means that all essential services should be started, but does not specify whether it should be started before or after it. References: systemd System and Service Manager; systemd.special(7) - Linux manual page

NEW QUESTION: 91
A Linux administrator needs to resolve a service that has failed to start. The administrator runs the following command:

```
ls -l startup file
```

The following output is returned

```
----------. root root 81k Sep 13 19:01 startupfile
```

Which of the following is MOST likely the issue?
A. The service does not have permissions to read write the startupfile.
B. The service startupfile size cannot be 81k.
C. The service startupfile cannot be owned by root.
D. The service startupfile should not be owned by the root group.
Answer: A (LEAVE A REPLY)
Explanation
The most likely issue is that the service does not have permissions to read or write the startupfile. The output of systemct1 status startup.service shows that the service has failed to start and the error message is
"Permission denied". The output of ls -l /etc/startupfile shows that the file has the permissions -rw-r--r--, which means that only the owner (root) can read and write the file, while the group (root) and others can only read the file. The service may not run as root and may need write access to the file. The administrator should change the permissions of the file by using the chmod command and grant write access to the group or others, or change the owner or group of the file by using the chown command and assign it to the user or group that runs the service. The other options are incorrect because they are not supported by the outputs. The file size, owner, and group are not the causes of the issue. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, pages 345-346.

NEW QUESTION: 92
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: (SHOW ANSWER)
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: 93
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 --state exited
C. docker rm $(docker ps -aq)
D. docker images prune *
Answer: (SHOW ANSWER)
Explanation
The command docker rm $(docker ps -aq) will allow the administrator to clean up the containers in an exited state. The docker command is a tool for managing Docker containers on Linux systems. Docker containers are isolated and lightweight environments that can run applications and services without affecting the host system.
Docker uses images to create containers, which are files that contain the code, libraries, dependencies, and configuration of the applications and services. The rm option removes one or more containers. The $(docker ps -aq) is a command substitution that executes

the command inside the parentheses and replaces it with the output. The docker ps -aq command lists all the containers, including the ones in an exited state, and shows only their IDs. The docker rm $(docker ps -aq) command will remove all the containers, including the ones in an exited state, by passing their IDs to the rm option. This will allow the administrator to clean up the containers in an exited state. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not exist (docker rm --all or docker rm --state exited) or do not remove the containers (docker images prune *). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 571.

NEW QUESTION: 94
After installing a new version of a package, a systems administrator notices a new version of the corresponding, service file was Installed In order to use the new version of the, service file, which of the following commands must be Issued FIRST?
A. systemctl status
B. systemctl stop
C. systemctl daemon-reload
D. systemctl reinstall
Answer: C (LEAVE A REPLY)

NEW QUESTION: 95
A Linux administrator is reviewing changes to a configuration file that includes the following section:

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


The Linux administrator is trying to select the appropriate syntax formatter to correct any issues with the configuration file. Which of the following should the syntax formatter support to meet this goal?
A. Markdown
B. XML
C. YAML
D. JSON
Answer: C (LEAVE A REPLY)
Explanation
The configuration file shown in the image is written in YAML format, so the syntax formatter should support YAML to correct any issues with the file. YAML stands for YAML Ain't Markup Language, and it is a human-readable data serialization language that uses indentation and colons to define key-value pairs. YAML supports various data types, such as scalars, sequences, mappings, anchors, aliases, and tags. The configuration file follows the rules and syntax of YAML, while the other options do not. Markdown is a lightweight markup language that uses plain text formatting to create rich text documents. XML is a markup language that uses tags to enclose elements and

attributes. JSON is a data interchange format that uses curly braces to enclose objects and square brackets to enclose arrays.
References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 21: Automating Tasks with Ansible, page 591.

NEW QUESTION: 96
A Linux administrator is trying to start the database service on a Linux server but is not able to run it. The administrator executes a few commands and receives the following output:

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

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

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


Which of the following should the administrator run to resolve this issue? (Select two).
A. systemctl unmask mariadb
B. journalctl -g mariadb
C. dnf reinstall mariadb
D. systemctl start mariadb
E. chkconfig mariadb on
F. service mariadb reload
Answer: A,D (LEAVE A REPLY)
Explanation
These commands will unmask the mariadb service, which is currently prevented from starting, and then start it normally. The other commands are either not relevant, not valid, or not sufficient for this task. For more information on how to manage masked services with systemctl, you can refer to the web search result 1.

NEW QUESTION: 97
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_nologin. so to the /etc/pam.d/sshd file.
D. Set PubKeyAuthentication to no in the /etc/ssh/ssh_config file.
Answer: B (LEAVE A REPLY)
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: 98
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 /.

Answer:

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: 99
A systems administrator is deploying three identical, cloud-based servers. The administrator is using the following code to complete the task:

```
resource "abc_instance" "ec2_instance" {

   ami                                = data.abc_ami.vendor-Linux-2.id
   associate_public_ip_address        = true
   count                              = 3
   instance_type                      = "instance_type"
   vpc_security_group_ids             = [abc.security_group.allow_ssh.
                                        id]
   key_name                           = abc_key_pair.key_pair.key_name

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

}
```


Which of the following technologies is the administrator using?
A. Puppet
B. Terraform
C. Chef
D. Ansible
Answer: B (LEAVE A REPLY)

NEW QUESTION: 100
A systems administrator pressed Ctrl+Z after starting a program using the command line, and the shell prompt was presented. In order to go back to the program, which of the following commands can the administrator use?
A. fg
B. su
C. bg
D. ed
Answer: A (LEAVE A REPLY)
Explanation
Ctrl+Z suspended the process, and "fg" will bring it back into the foreground of the shell A Comprehensive and Detailed Explanation: To go back to a program that was suspended by pressing Ctrl+Z in the command line, the command that can be used is fg. The fg command stands for foreground, and it resumes the job that is next in the queue and brings it to the foreground. Alternatively, if there are more than one suspended jobs, fg can be followed by a job number to resume a specific job. The other commands are incorrect because they either do not resume a suspended job, or they have different functions such as switching user (su), pushing a job to the background (bg), or editing a file (ed). References: CompTIA Linux+ Study Guide, Fourth Edition, page 181-182.

NEW QUESTION: 101
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. ssh -i "ABC-key.pem" root@l0.0.0.1
B. sftp rooteiO.0.0.1
C. scp "ABC-key.pem" root@10.0.0.1
D. sftp "ABC-key.pem" root@10.0.0.1
E. telnet 10.0.0.1 80
Answer: A (LEAVE A REPLY)

NEW QUESTION: 102
A systems administrator wants to back up the directory /data and all its contents to /backup/data on a remote server named remote. Which of the following commands will achieve the desired effect?
A. scp -p /data remote:/backup/data
B. ssh -i /remote:/backup/ /data
C. rsync -a /data remote:/backup/
D. cp -r /data /remote/backup/
Answer: C (LEAVE A REPLY)
Explanation
The command that will back up the directory /data and all its contents to /backup/data on a remote server named remote is rsync -a /data remote:/backup/. This command uses the rsync tool, which is a remote and local file synchronization tool. It uses an algorithm to minimize the amount of data copied by only moving the portions of files that have changed. The -a option stands for archive mode, which preserves the permissions, ownership, timestamps, and symbolic links of the files. The /data argument specifies the source directory to be backed up, and the remote:/backup/ argument specifies the destination directory on the remote server. The rsync tool will create a subdirectory named data under /backup/ on the remote server, and copy all the files and subdirectories from /data on the local server.

The other options are not correct commands for backing up a directory to a remote server. The scp -p /data remote:/backup/data command will copy the /data directory as a file named data under /backup/ on the remote server, not as a subdirectory with its contents. The -p option preserves the permissions and timestamps of the file, but not the ownership or symbolic links. The ssh -i /remote:/backup/ /data command will try to use /remote:/backup/ as an identity file for SSH authentication, which is not valid. The cp -r /data /remote/backup/ command will try to copy the /data directory to a local directory named /remote/backup/, not to a remote server. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks; rsync(1) - Linux manual page

NEW QUESTION: 103
A Linux user reported the following error after trying to connect to the system remotely:
ssh: connect to host 10.0.1.10 port 22: Resource temporarily unavailable The Linux systems administrator executed the following commands in the Linux system while trying to diagnose this issue:

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

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


Which of the following commands will resolve this issue?
A. firewall-cmd --zone=public --permanent --add-service=22
B. systemctl enable firewalld; systemctl restart firewalld
C. firewall-cmd --zone=public --permanent --add-service=ssh
D. firewall-cmd --zone=public --permanent --add-port=22/udp
Answer: (SHOW ANSWER)
The firewall-cmd --zone=public --permanent --add-service=ssh command will resolve the issue by allowing SSH connections on port 22 in the public zone of the firewalld service. This command will add the ssh service to the permanent configuration of the public zone, which means it will persist after a reboot or a reload of the firewalld service. The firewall-cmd --zone=public --permanent --add-service=22 command is invalid, as 22 is not a valid service name. The systemctl enable firewalld; systemctl restart firewalld command will enable and restart the firewalld service, but it will not change the firewall rules. The firewall-cmd --zone=public --permanent --add-port=22/udp

command will allow UDP traffic on port 22 in the public zone, but SSH uses TCP, not UDP. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 543.

NEW QUESTION: 104

A systems administrator is investigating why one of the servers has stopped connecting to the internet.

```
#curl http://google.com
curl: (6) Could not resolve host: google.com

#cat /etc/resolv.conf
search user.company.com company.com
#nameserver 10.10.10.10

#ip route
0.0.0.0/0 via 10.0.5.1 dev eth0 proto static metric 100
10.0.0.0/16 dev eth0 proto kernel scope link src 10.0.3.60 metric 101

#nmcli connection show
NAME                    UUID                                  TYPE          DEVICE
eth0                    ba4a3d30-efdc-4fa5-83d3-3721fd4aff75  ethernet      eth0
Wired connection 1      8d569d5a-22a2-356d-8532-9a2638f11b5a5  ethernet      --
```


Which of the following is causing the issue?

A. Wired connection 1 is offline.

B. The search entry in the /etc/resolv.conf file is incorrect.

C. No default route is defined.

D. The DNS address has been commented out in the configuration file.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 105

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

A. pam_login.so

B. pam_access.so

C. pam_logindef.so

D. pam_nologin.so

Answer: (SHOW ANSWER)

Explanation

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

NEW QUESTION: 106

Based on an organization's new cybersecurity policies, an administrator has been instructed to ensure that, by default, all new users and groups that are created fall within the specified values below.

```
# Min/max values for automatic uid selection in useradd
#
UID_MIN 1000
UID_MAX 60000
# Min/max values for automatic gid selection in groupadd
#
GID_MIN 1000
GID_MAX 60000
```


To which of the following configuration files will the required changes need to be made?

A. /etc/login.defs

B. /etc/default/useradd

C. /etc/profile

D. /etc/security/limits.conf

Answer: A (LEAVE A REPLY)

NEW QUESTION: 107

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

```
# df -h /ftpusers/

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

# df -i /ftpusers/

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


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

A. The users do not have the correct permissions to create files on the FTP server.

B. The ftpusers filesystem does not have enough space.

C. The inodes is at full capacity and would affect file creation for users.

D. ftpusers is mounted as read only.

Answer: C (LEAVE A REPLY)

Explanation

The cause of the issue based on the output above is C. The inodes is at full capacity and would affect file creation for users.

An inode is a data structure that stores information about a file or directory, such as its name, size, permissions, owner, timestamps, and location on the disk. Each file or directory has a unique inode number that identifies it. The number of inodes on a filesystem is fixed when the filesystem is created, and it determines how many files and directories can be created on that filesystem. If the inodes are exhausted, no new files or directories can be created, even if there is enough disk space available.

The output for the second command shows that the /ftpusers/ filesystem has 0% of inodes available, which means that all the inodes have been used up. This would prevent users from creating new files on the FTP server. The administrator should either delete some unused files or directories to free up some inodes, or resize the filesystem to increase the number of inodes.

The other options are incorrect because:

A: The users do not have the correct permissions to create files on the FTP server.

This is not true, because the output for the first command shows that the /ftpusers/ filesystem has 26% of disk space available, which means that there is enough space for users to create files. The permissions of the files and directories are not shown in the output, but they are not relevant to the issue of inode exhaustion.

B: The ftpusers filesystem does not have enough space.

This is not true, because the output for the first command shows that the /ftpusers/ filesystem has 26% of disk space available, which means that there is enough space for users to create files. The issue is not related to disk space, but to inode capacity.

D: ftpusers is mounted as read only.

This is not true, because the output for the first command does not show any indication that the /ftpusers/ filesystem is mounted as read only. If it was, it would have an (ro) flag next to the mounted on column. A read only filesystem would prevent users from creating or modifying files on the FTP server, but it would not affect the inode usage.

NEW QUESTION: 108

A Linux systems administrator is troubleshooting an I/O latency on a single CPU server. The administrator runs a top command and receives the following output:
%Cpu(s): 0.2 us, 33.1 sy, 0.0 ni, 0.0 id, 52.4 wa, 0.0 hi, 0.2 si, 0.0 st Which of the following is correct based on the output received from the executed command?
A. The server's CPU is taking too long to process users' requests.
B. The server's CPU shows a high idle-time value.
C. The server's CPU is spending too much time waiting for data inputs.
D. The server's CPU value for the time spent on system processes is low.
Answer: (SHOW ANSWER)
The server's CPU is spending too much time waiting for data inputs. This can be inferred from the output of the top command, which shows the percentage of CPU time spent in different states. The wa state stands for wait, and it indicates that the CPU is idle while waiting for an I/O operation to complete. In this case, the wa state is 52.4%, which means that more than half of the CPU time is wasted on waiting for data inputs. This can cause a high I/O latency and affect the performance of the server.
The other options are not correct based on the output received from the executed command. The server's CPU is not taking too long to process users' requests, because the us state, which stands for user, is only 0.2%, which means that the CPU is barely used by user processes. The server's CPU does not show a high idle-time value, because the id state, which stands for idle, is 0.0%, which means that the CPU is not idle at all. The server's CPU value for the time spent on system processes is not low, because the sy state, which stands for system, is 33.1%, which means that the CPU is heavily used by system processes.

NEW QUESTION: 109
A developer has been unable to remove a particular data folder that a team no longer uses. The developer escalated the issue to the systems administrator. The following output was received:

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


Which of the following commands can be used to resolve this issue?
A. chown -R data/
B. chattr -R -i data/
C. chmod -R 777 data/
D. chgrp -R 755 data/
Answer: B (LEAVE A REPLY)

NEW QUESTION: 110
A systems administrator installed a new software program on a Linux server. When the systems administrator tries to run the program, the following message appears on the screen.

```
Hardware virtualization support is not available on this system.
Either is not present or disabled in the system's BIOS
```

Which of the following commands will allow the systems administrator to check whether the system supports virtualization?
A. dmidecode -s system-version
B. lscpu
C. sysctl -a
D. cat /sys/device/system/cpu/possible
Answer: (SHOW ANSWER)
Explanation
The command that will allow the systems administrator to check whether the system supports virtualization is lscpu. This command will display information about the CPU architecture, such as the number of CPUs, cores, sockets, threads, model name, frequency, cache size, and flags. One of the flags is vmx (for Intel processors) or svm (for AMD processors), which indicates that the CPU supports hardware virtualization. If the flag is present, it means that the system supports virtualization. If the flag is absent, it means that the system does not support virtualization or that it is disabled in the BIOS settings.
The other options are not correct commands for checking whether the system supports virtualization. The dmidecode -s system-version command will display the version of the system, such as the product name or serial number, but not the CPU information. The sysctl -a command will display all the kernel parameters, but not the CPU flags. The cat /sys/devices/system/cpu/possible command will display the range of possible CPUs that can be online or offline, but not the CPU features. References: lscpu(1) - Linux manual page; How To Check If Virtualization is Enabled in Windows 10 / 11

NEW QUESTION: 111
A newly created container has been unable to start properly, and a Linux administrator is analyzing the cause of the failure. Which of the following will allow the administrator to determine the FIRST command that is executed inside the container right after it starts?
A. docker inspect <container_id>
B. docker info <container_id>
C. docker start <container_id>
D. docker export <container_id>
Answer: A (LEAVE A REPLY)

NEW QUESTION: 112
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 (LEAVE A REPLY)
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: 113
Several users reported that they were unable to write data to the /oracle1 directory. The following output has been provided:

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

Which of the following commands should the administrator use to diagnose the issue?
A. df -i /oracle1
B. fdisk -1 /dev/sdb1
C. lsblk /dev/sdb1
D. du -sh /oracle1
Answer: (SHOW ANSWER)
Explanation
The administrator should use the command df -i /oracle1 to diagnose the issue of users being unable to write data to the /oracle1 directory. This command will show the inode usage of the /oracle1 filesystem, which indicates how many files and directories can be created on it. If the inode usage is 100%, it means that no more files or directories can be added, even if there is still free space on the disk. The administrator can then delete some unnecessary files or directories, or increase the inode limit of the filesystem, to resolve the issue.
The other options are not correct commands for diagnosing this issue. The fdisk -l /dev/sdb1 command will show the partition table of /dev/sdb1, which is not relevant to the inode usage. The lsblk /dev/sdb1 command will show information about /dev/sdb1 as a block device, such as its size, mount point, and type, but not its inode usage. The du -sh /oracle1 command will show the disk usage of /oracle1 in human-readable format, but not its inode usage. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 7: Managing Disk Storage; How to Check Inode Usage in Linux - Fedingo

NEW QUESTION: 114
A Linux systems administrator needs to persistently enable IPv4 forwarding in one of the Linux systems. Which of the following commands can be used together to accomplish this task? (Choose two.)
A. echo "net.ipv4.ip_forward=1" >> /etc/sysctl.conf
B. sysctl net.ipv4.ip_forward
C. sysctl -p
D. echo 1 > /proc/sys/net/ipv4/ip_forward
E. echo "net.ipv6.conf.all.forwarding=l" >> /etc/sysctl.conf
F. sysctl -w net.ipv4.ip_forward=1
Answer: (SHOW ANSWER)

NEW QUESTION: 115
A Linux administrator has set up a new DNS forwarder and is configuring all internal servers to use the new forwarder to look up external DNS requests. The administrator needs to modify the firewall on the server for the DNS forwarder to allow the internal servers to communicate to it and make the changes persistent between server reboots. Which of the following commands should be run on the DNS forwarder server to accomplish this task?
A. ufw allow out dns
B. flrewall-cmd --zone-public --add-port-53/udp --permanent
C. iptables -A OUTPUT -p udp -ra udp -dport 53 -j ACCEPT
D. systemctl reload firewalld
Answer: B (LEAVE A REPLY)

NEW QUESTION: 116
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: (SHOW ANSWER)
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: 117
A Linux administrator wants to find out whether files from the wget package have been altered since they were installed. Which of the following commands will provide the correct information?
A. rpm -i wget
B. rpm -qf wget
C. rpm -F wget
D. rpm -V wget
Answer: (SHOW ANSWER)
The command that will provide the correct information about whether files from the wget package have been altered since they were installed is rpm -V wget. This command will use the rpm utility to verify an installed RPM package by comparing information about the installed files with information from the RPM database. The verification process can check various attributes of each file, such as size, mode, owner, group, checksum, capabilities, and so on. If any discrepancies are found, rpm will report them using a single letter code for each attribute.
The other options are not correct commands for verifying an installed RPM package. The rpm -i wget command is invalid because -i is used to install a package from a file, not to verify an installed package. The rpm -qf wget command will query which package owns wget as a file name or path name, but it will not verify its attributes. The rpm -F wget command will freshen (upgrade) an already installed package

with wget as a file name or path name, but it will not verify its attributes. Reference: rpm(8) - Linux manual page; Using RPM to Verify Installed Packages

NEW QUESTION: 118
A systems administrator is checking the system logs. The administrator wants to look at the last 20 lines of a log. Which of the following will execute the command?
A. tail -v 20
B. tail -n 20
C. tail -c 20
D. tail -l 20
Answer: B (LEAVE A REPLY)
Explanation
The command tail -n 20 will display the last 20 lines of a file. The -n option specifies the number of lines to show. This is the correct command to execute the task. The other options are incorrect because they either use the wrong options (-v, -c, or -l) or have the wrong arguments (20 instead of 20 filename). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, page 352.

NEW QUESTION: 119
A Linux administrator is troubleshooting a memory-related issue. Based on the output of the commands:

```
$ vmstat -s --unit M
968 M total memory
331 M used memory
482 M active memory
279 M inactive memory
 99 M free memory

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

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

Which of the following commands would address the issue?
A. top -p 8321
B. kill -9 8321
C. renice -10 8321

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

NEW QUESTION: 120
Joe, a user, is unable to log in to the Linux system Given the following output:

```
# grep joe /etc/passwd /etc/shadow
/etc/passwd:joe:x:1001:1001::/home/joe:/bin/nologin
/etc/shadow:joe:$6$3uOw6qWx9876jGhgKJsdfH987634534voj.:18883:0:99999:7:::
```

Which of the following command would resolve the issue?
A. usermod -s /bin/bash joe
B. pam_tally2 -u joe -r
C. passwd -u joe
D. chage -E 90 joe
Answer: B (LEAVE A REPLY)
Explanation
Based on the output of the image sent by the user, Joe is unable to log in to the Linux system because his account has been locked due to too many failed login attempts. The pam_tally2 -u joe -r command will resolve this issue by resetting Joe's failed login counter to zero and unlocking his account. This command uses the pam_tally2 module to manage user account locking based on login failures. The usermod -s /bin/bash joe command will change Joe's login shell to /bin/bash, but this will not unlock his account. The passwd -u joe command will unlock Joe's password if it has been locked by passwd -l joe, but this will not reset his failed login counter or unlock his account if it has been locked by pam_tally2. The chage -E 90 joe command will set Joe's account expiration date to 90 days from today, but this will not unlock his account or reset his failed login counter. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 537.

NEW QUESTION: 121
A systems administrator wants to list all local accounts in which the UID is greater than 500. Which of the following commands will give the correct output?
A. find /etc/passwd -size +500
B. cut -d: fl / etc/ passwd > 500
C. awk -F: '$3 > 500 {print $1}' /etc/passwd

D. sed '/UID/' /etc/passwd < 500
Answer: C (LEAVE A REPLY)
The correct command to list all local accounts in which the UID is greater than 500 is:
awk -F: '$3 > 500 {print $1}' /etc/passwd
This command uses awk to process the /etc/passwd file, which contains information about the local users on the system. The -F: option specifies that the fields are separated by colons. The $3 refers to the third field, which is the UID. The condition $3 > 500 filters out the users whose UID is greater than 500. The action {print $1} prints the first field, which is the username.
The other commands are incorrect because:
find /etc/passwd -size +500 will search for files that are larger than 500 blocks in size, not users with UID greater than 500.
cut -d: fl / etc/ passwd > 500 will cut the first field of the /etc/passwd file using colon as the delimiter, but it will not filter by UID or print only the usernames. The > 500 part will redirect the output to a file named 500, not compare with the UID.
sed '/UID/' /etc/passwd < 500 will use sed to edit the /etc/passwd file and replace any line that contains UID with 500, not list the users with UID greater than 500. The < 500 part will redirect the input from a file named 500, not compare with the UID.
Reference:
Linux List All Users In The System Command - nixCraft, section "List all users in Linux using /etc/passwd file".
Unix script getting users with UID bigger than 500 - Stack Overflow, section "Using awk".

NEW QUESTION: 122
A developer needs to launch an Nginx image container, name it Web001, and ex-pose port 8080 externally while mapping to port 80 inside the container. Which of the following commands will accomplish this task?
A. docker exec -it -p 8080: 80 --name Web001 nginx
B. docker load -it -p 8080:80 --name Web001 nginx
C. docker run -it -P 8080:80 --name Web001 nginx
D. docker pull -it -p 8080:80 -name Web00l nginx
Answer: C (LEAVE A REPLY)
Explanation
To launch an Nginx image container, name it Web001, and expose port 8080 externally while mapping to port
80 inside the container, the administrator can use the command docker run -it -p 8080:80 --name Web001 nginx . This will create and start a new container from the Nginx image, assign it a name of Web001, and map port 8080 on the host to port 80 on the container. The other commands are not valid or do not meet the requirements. References:
[CompTIA Linux+ Study Guide], Chapter 11: Working with Containers, Section: Running Containers with Docker
[How to Run Docker Containers]

NEW QUESTION: 123

The security team has identified a web service that is running with elevated privileges A Linux administrator is working to change the systemd service file to meet security compliance standards. Given the following output:

```
[Unit]
Description=CompTIA server daemon
Documentation=man:webserver(8) man:webserver_config(5)
After=network.target

[Service]
Type=notify
EnvironmentFile=/etc/webserver/config
ExecStart=/usr/sbin/webserver -D $OPTIONS
ExecReload=/bin/kill -HUP $MAINPID
KillMode=process
Restart=on-failure
RestartSec=42s

[Install]
WantedBy=multi-user.target
```


Which of the following remediation steps will prevent the web service from running as a privileged user?

A. Removing the ExecStarWusr/sbin/webserver -D SOPTIONS from the service file

B. Updating the Environment File line in the [Service] section to/home/webservice/config

C. Adding the User-webservice to the [Service] section of the service file

D. Changing the:nulti-user.target in the [Install] section to basic.target

Answer: C (LEAVE A REPLY)

Explanation

The remediation step that will prevent the web service from running as a privileged user is adding the User=webservice to the [Service] section of the service file. The service file is a configuration file that defines the properties and behavior of a systemd service. The systemd is a system and service manager that controls the startup and operation of Linux systems. The service file contains various sections and options that specify how the service should be started, stopped, and managed. The [Service] section defines how the service should be executed and what commands should be run. The User option specifies the user name or ID that the service should run as. The webservice is the name of the user that the administrator wants to run the web service as.

The administrator should add the User=webservice to the [Service] section of the service file, which will prevent the web service from running as a privileged user, such as root, and improve the security of the system. This is the correct remediation step to use to prevent the web service from running as a privileged user. The other options are incorrect because they either do not change the user that the service runs as (removing the ExecStart=/usr/sbin/webserver -D OPTIONS from the service file or updating the EnvironmentFile line in the [Service] section to /home/webservice/config) or do not affect the user that the service runs as (changing the multi-user.target in the [Install] section to basic.target). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing System Services, page 458.

NEW QUESTION: 124

A Linux engineer is setting the sticky bit on a directory called devops with 755 file permission. Which of the following commands will accomplish this task?

A. chown -s 755 devops

B. chown 1755 devops

C. chmod -s 755 devops

D. chmod 1755 devops

Answer: D (LEAVE A REPLY)

Explanation

The command that will set the sticky bit on a directory called devops with 755 file permission is chmod 1755 devops. This command will use chmod to change the mode of the directory devops to 1755, which means that the owner has read, write, and execute permissions (7), the group has read and execute permissions (5), and others have read and execute permissions (5). The first digit 1 indicates that the sticky bit is set on the directory, which is a special permission that prevents users from deleting or renaming files in the directory that they do not own.

The other options are not correct commands for setting the sticky bit on a directory. The chown -s 755 devops command is invalid because chown is used to change the owner and group of files or directories, not their permissions. The -s option for chown is used to remove a symbolic link, not to set the sticky bit. The chown

1755 devops command is also invalid because chown does not accept numeric arguments for changing permissions. The chmod -s 755 devops command will remove the sticky bit from the directory devops, not set it. References: chmod(1) - Linux manual page; How to Use SUID, SGID, and Sticky Bits on Linux

NEW QUESTION: 125

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	9435	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 (LEAVE A REPLY)

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: 126
A Linux administrator is troubleshooting the root cause of a high CPU load and average.

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

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

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


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

NEW QUESTION: 127
A Linux system is failing to start due to issues with several critical system processes. Which of the following options can be used to boot the system into the single user mode? (Choose two.)
A. Execute the following command from the GRUB rescue shell: mount -o remount, ro/sysroot.
B. Interrupt the boot process in the GRUB menu and add systemd.unit=single in the kernel line.
C. Interrupt the boot process in the GRUB menu and add systemd.unit=rescue.target in the kernel line.
D. Interrupt the boot process in the GRUB menu and add single=user in the kernel line.

E. Interrupt the boot process in the GRUB menu and add init=/bin/bash in the kernel line.
F. Interrupt the boot process in the GRUB menu and add systemd.unit=single.target in the kernel line.
Answer: (SHOW ANSWER)
The administrator can use the following two options to boot the system into the single user mode:
Interrupt the boot process in the GRUB menu and add systemd.unit=rescue.target in the kernel line. This option will boot the system into the rescue mode, which is a minimal environment that allows the administrator to perform basic tasks such as repairing the system. The GRUB menu is a screen that appears when the system is powered on and allows the administrator to choose which kernel or operating system to boot. The kernel line is a line that specifies the parameters for the kernel, such as the root device, the init system, and the boot options. The administrator can interrupt the boot process by pressing the e key in the GRUB menu and edit the kernel line by adding systemd.unit=rescue.target at the end. This option will tell the system to use the rescue target, which is a unit that defines the state of the system in the rescue mode. The administrator can then press Ctrl+X to boot the system with the modified kernel line. This option will boot the system into the single user mode and allow the administrator to troubleshoot the issues.
Interrupt the boot process in the GRUB menu and add systemd.unit=single.target in the kernel line. This option will boot the system into the single user mode, which is a mode that allows the administrator to log in as the root user and perform maintenance tasks. The GRUB menu and the kernel line are the same as the previous option. The administrator can interrupt the boot process by pressing the e key in the GRUB menu and edit the kernel line by adding systemd.unit=single.target at the end. This option will tell the system to use the single target, which is a unit that defines the state of the system in the single user mode. The administrator can then press Ctrl+X to boot the system with the modified kernel line. This option will boot the system into the single user mode and allow the administrator to troubleshoot the issues.
The other options are incorrect because they either do not boot the system into the single user mode (execute the following command from the GRUB rescue shell: mount -o remount, ro/sysroot or interrupt the boot process in the GRUB menu and add systemd.unit=single in the kernel line) or do not use the correct syntax (interrupt the boot process in the GRUB menu and add single=user in the kernel line or interrupt the boot process in the GRUB menu and add init=/bin/bash in the kernel line). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing the Linux Boot Process, pages 267-268.

NEW QUESTION: 128
Users in the human resources department are trying to access files in a newly created directory. Which of the following commands will allow the users access to the files?
A. chattr
B. chgrp
C. chage
D. chcon
Answer: B (LEAVE A REPLY)
Explanation
The chgrp command is used to change the group ownership of files and directories. By using this command, the administrator can assign the files in the newly created directory to the human resources group, which will allow the users in that group to access them. The other commands are not relevant for this task. For example:
chattr is used to change the file attributes, such as making them immutable or append-only1.
chage is used to change the password expiration information for a user account2.
chcon is used to change the security context of files and directories, which is related to SELinux3.
References:

The CompTIA Linux+ Certification Exam Objectives mention that the candidate should be able to
"manage file and directory ownership and permissions" as part of the Hardware and System Configuration domain4.
The web search result 2 explains how to use the chgrp command with examples.
The web search result 3 compares the chmod and chgrp commands and their effects on file permissions.

NEW QUESTION: 129
A Linux engineer has been notified about the possible deletion of logs from the file /opt/app/logs. The engineer needs to ensure the log file can only be written into without removing previous entries.

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

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

NEW QUESTION: 130
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	[illegible]	[illegible]	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 (LEAVE A REPLY)
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: 131
A Linux administrator rebooted a server. Users then reported some of their files were missing. After doing some troubleshooting, the administrator found one of the filesystems was missing. The filesystem was not listed in /etc/f stab and might have been mounted manually by someone prior to reboot. Which of the following would prevent this issue from reoccurring in the future?
A. Sync the mount units.
B. Create a mount unit and enable it to be started at boot.
C. Mount the filesystem manually.
D. Remount all the missing filesystems
Answer: (SHOW ANSWER)

NEW QUESTION: 132
Which of the following data structures is written in JSON?

A.
```
---
name: user1
position: DevOps
floor: 3
```

B.
```
<table>
<tbody><tr>
<td>user1</td>
<td>DevOps</td>
<td>3</td>
</tr>
</tbody></table>
```

C.
```
<root>
   <floor>3</floor>
   <name>user1</name>
   <position>DevOps</position>
</root>
```

D.
```
{
   "name": "user1",
   "job": "DevOps",
   "floor": 3
}
```

Answer: C (LEAVE A REPLY)
Explanation

Option C is the only data structure that is written in JSON format. JSON stands for JavaScript Object Notation, and it is a lightweight and human-readable data interchange format. JSON uses curly braces to enclose objects, which consist of key-value pairs separated by commas. JSON uses square brackets to enclose arrays, which consist of values separated by commas. JSON supports six data types: strings, numbers, booleans, null, objects, and arrays. Option C follows these rules and syntax of JSON, while the other options do not. Option A is written inXML format, which uses tags to enclose elements and attributes. Option B is written in YAML format, which uses indentation and colons to define key-value pairs. Option D is written in INI format, which uses sections and equal signs to define key-value pairs. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 21: Automating Tasks with Ansible, page 591.

NEW QUESTION: 133
A Linux systems administrator needs to copy files and directories from Server A to Server B. Which of the following commands can be used for this purpose? (Select TWO)
A. reposync
B. cp
C. scp
D. rsync
E. ssh
F. rsyslog
Answer: C,D (LEAVE A REPLY)

NEW QUESTION: 134
After starting an Apache web server, the administrator receives the following error:
Apr 23 localhost.localdomain httpd 4618] : (98) Address already in use: AH00072: make_sock: could not bind to address [: :]80 Which of the following commands should the administrator use to further trou-bleshoot this issue?
A. Ss
B. Ip
C. Dig
D. Nc
Answer: (SHOW ANSWER)
Explanation
The ss command is used to display information about socket connections, such as the port number, state, and process ID. The error message indicates that the port 80 is already in use by another process, which prevents the Apache web server from binding to it. By using the ss command with the -l and -n options, the administrator can list all the listening sockets and their port numbers in numeric form, and identify which process is using the port 80. For example: ss -ln | grep :80. The ip, dig, and nc commands are not relevant for this issue, as they are used for different purposes, such as configuring network interfaces, querying DNS records, and testing network connectivity.

NEW QUESTION: 135
A Linux administrator would like to use systemd to schedule a job to run every two hours. The administrator creates timer and service definitions and restarts the server to load these new configurations. After the restart, the administrator checks the log file and notices that the job is only running daily. Which of the following is MOST likely causing the issue?
A. The checkdiskspace.service is not running.
B. The checkdiskspace.service needs to be enabled.

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

NEW QUESTION: 136
Based on an organization's new cybersecurity policies, an administrator has been instructed to ensure that, by default, all new users and groups that are created fall within the specified values below.

```
# Min/max values for automatic uid selection in useradd
#
UID_MIN 1000
UID_MAX 60000
# Min/max values for automatic gid selection in groupadd
#
GID_MIN 1000
GID_MAX 60000
```


To which of the following configuration files will the required changes need to be made?
A. /etc/login.defs
B. /etc/security/limits.conf
C. /etc/default/useradd
D. /etc/profile
Answer: (SHOW ANSWER)
Explanation
The required changes need to be made to the /etc/login.defs configuration file. The /etc/login.defs file defines the default values for user and group IDs, passwords, shells, and other parameters for user and group creation.
The file contains the directives UID_MIN, UID_MAX, GID_MIN, and GID_MAX, which set the minimum and maximum values for automatic user and group ID selection. The administrator can edit this file and change the values to match the organization's new cybersecurity policies. This is the correct file to modify to accomplish the task. The other options are incorrect because they either do not affect the user

and group IDs (/etc/security/limits.conf or /etc/profile) or do not set the default values (/etc/default/useradd). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15:
Managing Users and Groups, page 463.

NEW QUESTION: 137
A Linux administrator is troubleshooting a memory-related issue. Based on the output of the commands:

```
$ vmstat -s --unit M

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

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

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

Which of the following commands would address the issue?
A. top -p 8321
B. kill -9 8321
C. renice -10 8321
D. free 8321
Answer: (SHOW ANSWER)
Explanation
The command that would address the memory-related issue is kill -9 8321. This command will send a SIGKILL signal to the process with the PID 8321, which is the mysqld process that is using 99.7% of the available memory according to the top output. The SIGKILL signal will terminate the process immediately and free up the memory it was using. However, this command should be used with caution as it may cause data loss or corruption if the process was performing some critical operations.

The other options are not correct commands for addressing the memory-related issue. The top -p 8321 command will only display information about the process with the PID 8321, but will not kill it or reduce its memory usage. The renice -10 8321 command will change the priority (niceness) of the process with the PID
8321 to -10, which means it will have a higher scheduling priority, but this will not affect its memory consumption. The free 8321 command is invalid because free does not take a PID as an argument; free only displays information about the total, used, and free memory in the system. References: How to troubleshoot Linux server memory issues; kill(1) - Linux manual page

NEW QUESTION: 138
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 (LEAVE A REPLY)
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: 139
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 (LEAVE A REPLY)
git clone is the best option if its a new project. git pull would work best if it was an ongoing project that had multiple team members that had pushed their updates to the main repository

NEW QUESTION: 140
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: (SHOW ANSWER)
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: 141
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: (SHOW ANSWER)
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: 142
A Linux administrator created a new file system. Which of the following files must be updated to ensure the filesystem mounts at boot time?
A. /etc/sysctl
B. /etc/filesystems
C. /etc/fstab
D. /etc/nfsmount.conf
Answer: C (LEAVE A REPLY)
The file that must be updated to ensure the filesystem mounts at boot time is /etc/fstab. This file contains information about the filesystems that are mounted automatically by the mount -a command, which is usually invoked during the system startup. The /etc/fstab file has six fields for each filesystem: device name, mount point, filesystem type, mount options, dump frequency, and pass number. To add a new filesystem to the /etc/fstab file, you need to specify these fields correctly and make sure the mount point directory exists.
The other options are not correct files for controlling persistent mount points of filesystems. The /etc/sysctl file is used to configure kernel parameters at runtime. The /etc/filesystems file is used to specify the order of filesystem types used by mount when no filesystem type is given. The /etc/nfsmount.conf file is used to set options for mounting NFS filesystems. Reference: Persistently mounting file systems; fstab(5) - Linux manual page

NEW QUESTION: 143
Users have been unable to reach www.comptia.org from a Linux server. A systems administrator is troubleshooting the issue and does the following:

Output 1:

```
2: eth0: <BROADCAST,MULTICAST,UP, LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
   link/ether ac:11:22:33:44:cd brd ff:ff:ff:ff:ff:ff
   inet 192.168.168.10/24 brd 192.168.169.255 scope global dynamic noprefixroute eth0
      valid_lft 8097sec preferred_lft 8097sec
   inet fe80::4daf:8c7c:a6ff:2771/64 scope link noprefixroute
      valid_lft forever preferred_lft forever
```

Output 2:

```
nameserver 192.168.168.53
```

Output 3:

```
FING 192.168.168.53 (192.168.168.53) 56(84) bytes of data.
64 bytes from 192.168.168.53: icmp_seq=1 ttl=64 time=2.85 ms

--- 192.168.168.53 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 2.847/2.847/2.847/0.000 ms
```

Output 4:

```
192.168.168.0/24 dev eth0 proto kernel scope link src 192.168.168.10 metric 600
```

Output 5:

```
...
;; QUESTION SECTION:
;www.comptia.org. IN A

;; ANSWER SECTION:
. 0 CLASS4096 OPT 10 8 LgmNvk0AazU=

;; ADDITIONAL SECTION:
www.comptia.org. 3385 IN A 23.96.239.26
...
```

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

A. The name www.comptia.org does not point to a valid IP address.

B. The server 192.168.168.53 is unreachable.

C. No default route is set on the server.

D. The network interface eth0 is disconnected.

Answer: B (LEAVE A REPLY)

Explanation

The issue is caused by the server 192.168.168.53 being unreachable. This server is the DNS server configured in the /etc/resolv.conf file, which is used to resolve domain names to IP addresses. The ping command shows that the server cannot be reached, and the nslookup command shows that the name www.comptia.org cannot be resolved using this server. The other options are incorrect because:

The name www.comptia.org does point to a valid IP address, as shown by the nslookup command using another DNS server (8.8.8.8).

The default route is set on the server, as shown by the ip route command, which shows a default gateway of 192.168.168.1.

The network interface eth0 is connected, as shown by the ip link command, which shows a state of UP for eth0. References: CompTIA Linux+ Study Guide, Fourth Edition, page 457-458, 461-462.

NEW QUESTION: 144

A cloud engineer needs to launch a container named web-01 in background mode. Which of the following commands will accomplish this task"

A. docker ps -a --name web-01 httpd
B. docker load --name web-01 httpd
C. docker builder -f -name web-01 httpd
D. docker run -d --name web-01 httpd
Answer: D (LEAVE A REPLY)

NEW QUESTION: 145
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	9663	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. kill -9 "upload*.sh"
B. skill -9 "upload*.sh"
C. killall -9 -upload*.sh"
D. pkill -9 -f "upload*.sh"
Answer: C (LEAVE A REPLY)

NEW QUESTION: 146
A systems administrator is deploying three identical, cloud-based servers. The administrator is using the following code to complete the task:

```
resource "abc_instance" "ec2_instance" {

   ami                                  = data.abc_ami.vendor-Linux-2.id
   associate_public_ip_address          = true
   count                                = 3
   instance_type                        = "instance_type"
   vpc_security_group_ids               = [abc.security_group.allow_ssh.
                                          id]
   key_name                             = abc_key_pair.key_pair.key_name

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

}
```


Which of the following technologies is the administrator using?

A. Ansible

B. Puppet

C. Chef

D. Terraform

Answer: D (LEAVE A REPLY)

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

NEW QUESTION: 147

A Linux systems administrator needs to copy files and directories from Server A to Server B. Which of the following commands can be used for this purpose? (Select TWO)

A. rsyslog

B. cp

C. rsync

D. reposync

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

NEW QUESTION: 148
A DevOps engineer is working on a local copy of a Git repository. The engineer would like to switch from the main branch to the staging branch but notices the staging branch does not exist. Which of the following Git commands should the engineer use to perform this task?
A. git branch -m staging
B. git commit -m staging
C. git status -b staging
D. git checkout -b staging
Answer: D (LEAVE A REPLY)
Explanation
The correct answer is D. git checkout -b staging
This command will create a new branch named staging and switch to it. The git checkout command is used to switch between branches or restore files from a specific branch. The -b option is used to create a new branch if it does not exist. For example, git checkout -b staging will create and switch to the staging branch.
The other options are incorrect because:
A: git branch -m staging
This command will rename the current branch to staging, not switch to it. The git branch command is used to list, create, or delete branches. The -m option is used to rename a branch. For example, git branch -m staging will rename the current branch to staging.
B: git commit -m staging
This command will commit the changes in the working tree to the current branch with a message of staging, not switch to it. The git commit command is used to record changes to the repository. The -m option is used to specify a commit message. For example, git commit -m staging will commit the changes with a message of staging.
C: git status -b staging
This command will show the status of the working tree and the current branch, not switch to it. The git status command is used to show the state of the working tree and the staged changes. The -b option is used to show the name of the current branch. However, this option does not take an argument, so specifying staging after it will cause an error.
References:
Git - git-checkout Documentation
Git Tutorial: Create a New Branch With Git Checkout
Git Branching - Basic Branching and Merging

NEW QUESTION: 149

As a Systems Administrator, to reduce disk space, you were tasked to create a shell script that does the following:

Add relevant content to /tmp/script.sh, so that it finds and compresses rotated files in /var/log without recursion.

INSTRUCTIONS

Fill the blanks to build a script that performs the actual compression of rotated log files.

If at any time you would like to bring back the initial state of the simulation, please click the Reset All button.

Answer:

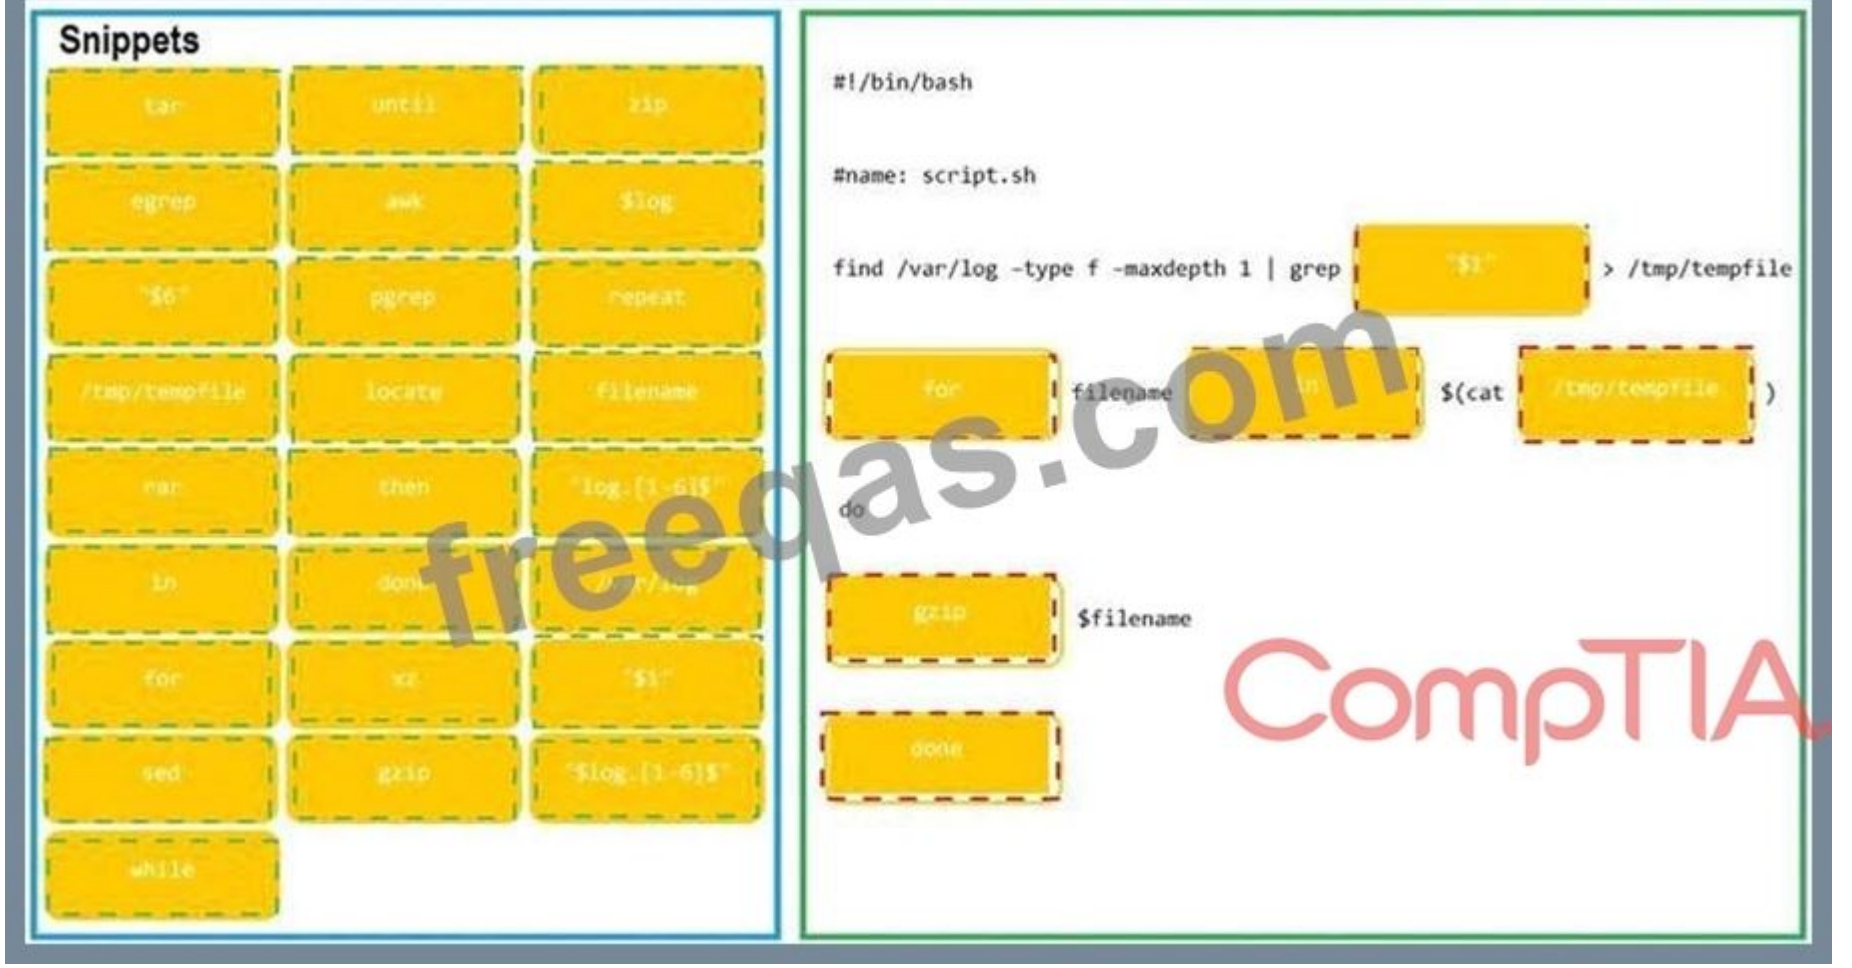
Snippets
tar
until
zip
egrep
awk
$log
"$6"
pgrep
repeat
/tmp/tempfile
locate
filename
rar
then
"log.[1-6]$"
in
done
for
xz
"$1"
sed
gzip
"$log.[1-6]$"
while
#!/bin/bash
#name: script.sh
find /var/log -type f -maxdepth 1 | grep "$1" > /tmp/tempfile
for filename in $(cat /tmp/tempfile)
do
gzip $filename
done
freeqas.com
CompTIA

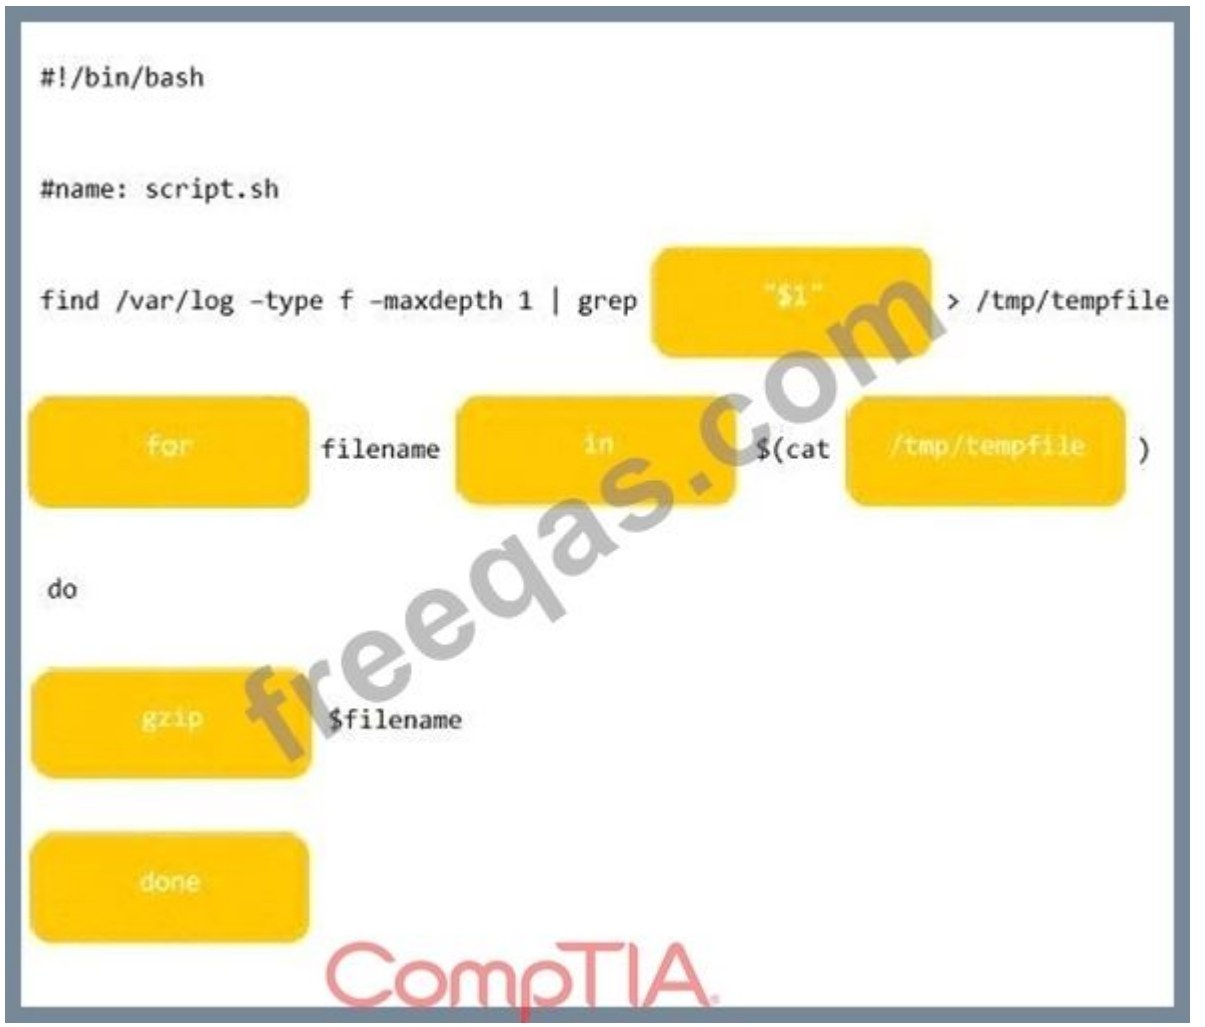

NEW QUESTION: 150
A new file was added to a main Git repository. An administrator wants to synchronize a local copy with the contents of the main repository. Which of the following commands should the administrator use for this task?
A. git status
B. git push
C. git reflog
D. git pull
Answer: (SHOW ANSWER)

NEW QUESTION: 151
A Linux administrator is troubleshooting an issue in which an application service failed to start on a Linux server. The administrator runs a few commands and gets the following outputs:

```
Output 1:

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

Output 2:

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

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

A. Enable the logsearch.service and restart the service.

B. Increase the TimeoutStartUSec configuration for the logsearch.sevice.

C. Update the OnCalendar configuration to schedule the start of the logsearch.service.

D. Update the KillSignal configuration for the logsearch.service to use TERM.

Answer: (SHOW ANSWER)

Explanation

The administrator should increase the TimeoutStartUSec configuration for the logsearch.service to resolve the issue. The output of systemct1 status logsearch.service shows that the service failed to start due to a timeout.

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

NEW QUESTION: 152

Users have been unable to reach www.comptia.org from a Linux server. A systems administrator is troubleshooting the issue and does the following:

Output 1:

```
2: eth0: <BROADCAST,MULTICAST,UP, LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
   link/ether ac:11:22:33:44:cd brd ff:ff:ff:ff:ff:ff
   inet 192.168.168.10/24 brd 192.168.169.255 scope global dynamic noprefixroute eth0
      valid_lft 8097sec preferred_lft 8097sec
   inet fe80::4daf:8c7c:a6ff:2771/64 scope link noprefixroute
      valid_lft forever preferred_lft forever
```

Output 2:

```
nameserver 192.168.168.53
```

Output 3:

```
FING 192.168.168.53 (192.168.168.53) 56(84) bytes of data.
64 bytes from 192.168.168.53: icmp_seq=1 ttl=64 time=2.85 ms

--- 192.168.168.53 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 2.847/2.847/2.847/0.000 ms
```

Output 4:

```
192.168.168.0/24 dev eth0 proto kernel scope link src 192.168.168.10 metric 600
```

Output 5:

```
...
;; QUESTION SECTION:
;www.comptia.org. IN A

;; ANSWER SECTION:
. 0 CLASS4096 OPT 10 8 LgmNvk0AazU=

;; ADDITIONAL SECTION:
www.comptia.org. 3385 IN A 23.96.239.26
...
```

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

A. The server 192.168.168.53 is unreachable.

B. No default route is set on the server.

C. The name www.comptia.org does not point to a valid IP address.

D. The network interface eth0 is disconnected.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 153

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 (LEAVE A REPLY)

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: 154

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 (LEAVE A REPLY)

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: 155

A user is asking the systems administrator for assistance with writing a script to verify whether a file exists. Given the following:

```
#1/bin/bash
filename=$1
<CONDITIONAL>
echo "File exists"
else
echo "File does not exist"
fi
```


Which of the following commands should replace the <CONDITIONAL> string?

A. if [-f "$filename"]; while

B. if [-f "$filename"]; then

C. if [-f "$filename"] then

D. if [-d "$filename"]; then

Answer: B (LEAVE A REPLY)

NEW QUESTION: 156

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 (LEAVE A REPLY)

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: 157

A Linux administrator needs to expand a volume group using a new disk. Which of the following options presents the correct sequence of commands to accomplish the task?

A. partprobe

vgcreate
lvextend
B. lvcreate
fdisk
partprobe
C. fdisk
partprobe
mkfs
D. fdisk
pvcreate
vgextend
Answer: D (LEAVE A REPLY)
The correct sequence of commands to expand a volume group using a new disk is fdisk, pvcreate, vgextend. The fdisk command can be used to create a partition on the new disk with the type 8e (Linux LVM). The pvcreate command can be used to initialize the partition as a physical volume for LVM. The vgextend command can be used to add the physical volume to an existing volume group. The partprobe command can be used to inform the kernel about partition table changes, but it is not necessary in this case. The vgcreate command can be used to create a new volume group, not expand an existing one. The lvextend command can be used to extend a logical volume, not a volume group. The lvcreate command can be used to create a new logical volume, not expand a volume group. The mkfs command can be used to create a filesystem on a partition or a logical volume, not expand a volume group. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Disk Storage, pages 462-463.

NEW QUESTION: 158
Which of the following can be used as a secure way to access a remote terminal?
A. TFTP
B. SSH
C. SCP
D. SFTP
Answer: B (LEAVE A REPLY)
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: 159

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

A. chmod 775

B. umask. 002

C. chactr -Rv

D. chown -cf

Answer: B (LEAVE A REPLY)

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

NEW QUESTION: 160

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. Whileworking 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 (LEAVE A REPLY)

Explanation

The command restorecon -R -v /var/www/html will best resolve the issue. The issue is caused by the incorrect SELinux context of the web server files under the /var/www/html directory. The output of ls -Z
/var/www/html shows that the files have the type user_home_t, which is not allowed for web content. The command restorecon restores the default SELinux context of files based on the policy rules. The options -R and -v are used to apply the command recursively and verbosely. This command will change the type of the files to httpd_sys_content_t, which is the correct type for web content. This will allow

the web server to access the files and serve the pages to the browser. The other options are incorrect because they either disable SELinux entirely (sed -i 's/SELINUX=enforcing/SELINUX=disabled/'
/etc/selinux/config or setenforce 0), which is not a good security practice, or enable an unnecessary boolean (setsebool -P httpd_can_network_connect_db on), which is not related to the issue. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Securing Linux Systems, page 535.

NEW QUESTION: 161
A Linux administrator needs to obtain a list of all volumes that are part of a volume group. Which of the following commands should the administrator use to accomplish this task?
A. pvs
B. vgs
C. fdisk -1
D. lvs
Answer: D (LEAVE A REPLY)

NEW QUESTION: 162
A systems administrator is implementing a new service task with systems at startup and needs to execute a script entitled test.sh with the following content:

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

The administrator tries to run the script after making it executable with chmod +x; however, the script will not run. Which of the following should the administrator do to address this issue? (Choose two.)
A. Create a unit file for the new service in /etc/init.d with the name helpme.service in the location.
B. Create a unit file for the new service in /etc/systemd/system/ with the name helpme.service in the location.
C. Add #!//bin/bash to the top of the script.
D. Shut down the computer to enable the new service.
E. Restart the computer to enable the new service.
F. Add #!/bin/bash to the bottom of the script.
Answer: (SHOW ANSWER)

NEW QUESTION: 163
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 (LEAVE A REPLY)

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: 164

A systems administrator is investigating an issue in which one of the servers is not booting up properly. The journalctl entries show the following:

```
Sep 16 20:30:43 server kernel: acpi PNP0A03:00: _OSC failed (AE_NOT_FOUND);
-- Subject: Unit dev-mapper-centos\x2dapp.device has failed
Sep 16 20:32:15 server systemd[1]: Dependency failed for /opt/app
-- Subject: Unit opt-app.mount has failed
-- Unit opt-app.mount has failed
Sep 16 20:32:15 server systemd[1]: Dependency failed for Local File Systems.
-- Subject: Unit local-fs.target has failed
-- Unit local-fs.target has failed.
Sep 16 20:32:15 server systemd[1]: Dependency failed for Relabel all filesystem, if necessary.
-- Subject: Unit rhel-autorelabel.service has failed
-- Unit rhel-autorelabel.service has failed.
```

Which of the following will allow the administrator to boot the Linux system to normal mode quickly?

A. Comment out the /opt/app filesystem in /etc/fstab and reboot.

B. Reformat the /opt/app filesystem and reboot.

C. Perform filesystem checks on local filesystems and reboot.

D. Trigger a filesystem relabel and reboot.

Answer: (SHOW ANSWER)

Explanation

The fastest way to boot the Linux system to normal mode is to comment out the /opt/app filesystem in /etc/fstab and reboot. This will prevent the system from trying to mount the /opt/app filesystem at boot time, which causes an error because the filesystem does not exist or is corrupted. Commenting out a line in /etc/fstab can be done by adding a # symbol at the beginning of the line. Rebooting the system will apply the changes and allow the system to boot normally. Reformatting the /opt/app filesystem will not help to boot the system, as it will erase any data on the filesystem and require manual intervention to create a new filesystem.

Performing filesystem checks on local filesystems will not help to boot the system, as it will not fix the missing or corrupted /opt/app filesystem. Triggering a filesystem relabel will not help to boot the system, as it will only change the security context of files and directories according to SELinux policy. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Disk Storage, page 456.

NEW QUESTION: 165

A developer wants to ensure that all files and folders created inside a shared folder named /GroupOODEV inherit the group name of the parent folder. Which of the following commands will help achieve this goal?

A. chmod g+X / GroupOODEV/

B. chmod g+W / GroupOODEV/
C. chmod g+r / GroupOODEV/
D. chmod g+s / GroupOODEV/
Answer: (SHOW ANSWER)
Explanation
The chmod command is used to change the permissions of files and directories on Linux systems.
The g+s option sets the setgid bit on a directory, which means that all files and folders created inside that directory will inherit the group name of the parent directory. This command can help the developer ensure that all files and folders created inside the /GroupOODEV directory have the same group name as
/GroupOODEV. References: [How to Use chmod Command in Linux with Examples]

NEW QUESTION: 166
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. Too many users are currently logged in to the system
C. The system has been running for over a year and requires a reboot.
D. The system requires more memory
Answer: C (LEAVE A REPLY)

NEW QUESTION: 167
A Linux administrator has set up a new DNS forwarder and is configuring all internal servers to use the new forwarder to look up external DNS requests. The administrator needs to modify the firewall on the server for the DNS forwarder to allow the internal servers to communicate to it and make the changes persistent between server reboots. Which of the following commands should be run on the DNS forwarder server to accomplish this task?
A. ufw allow out dns
B. systemct1 reload firewalld
C. iptables -A OUTPUT -p udp -ra udp -dport 53 -j ACCEPT
D. flrewall-cmd --zone-public --add-port-53/udp --permanent
Answer: (SHOW ANSWER)

Explanation
The command that should be run on the DNS forwarder server to accomplish the task is firewall-cmd --zone=public --add-port=53/udp --permanent. The firewall-cmd command is a tool for managing firewalld, which is a firewall service that provides dynamic and persistent network security on Linux systems. The firewalld uses zones and services to define the rules and policies for the network traffic. The zones are logical groups of network interfaces and sources that have the same level of trust and security. The services are predefined sets of ports and protocols that are associated with certain applications or functions.
The --zone=public option specifies the zone name that the rule applies to. The public zone is the default zone that represents the untrusted network, such as the internet. The --add-port=53/udp option adds a port and protocol to the zone. The 53 is the port number that is used by the DNS service. The udp is the protocol that is used by the DNS service. The --permanent option makes the change persistent across reboots. The command firewall-cmd --zone=public --add-port=53/udp --permanent will modify the firewall on the server for the DNS forwarder to allow the internal servers to communicate to it and make the changes persistent between server reboots. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not modify the firewall on the server for the DNS forwarder (ufw allow out dns or systemct1 reload firewalld) or do not use the correct syntax for the command (iptables -A OUTPUT -p udp -ra udp -dport 53 -j ACCEPT instead of iptables -A OUTPUT -p udp -ra udp --dport 53 -j ACCEPT). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 392.

NEW QUESTION: 168
A Linux administrator rebooted a server. Users then reported some of their files were missing. After doing some troubleshooting, the administrator found one of the filesystems was missing. The filesystem was not listed in /etc/f stab and might have been mounted manually by someone prior to reboot. Which of the following would prevent this issue from reoccurring in the future?
A. Sync the mount units.
B. Mount the filesystem manually.
C. Create a mount unit and enable it to be started at boot.
D. Remount all the missing filesystems
Answer: C (LEAVE A REPLY)
Explanation
The best way to prevent this issue from reoccurring in the future is to create a mount unit and enable it to be started at boot. A mount unit is a systemd unit that defines how and where a filesystem should be mounted. By creating a mount unit for the missing filesystem and enabling it with systemctl enable, the administrator can ensure that the filesystem will be automatically mounted at boot time, regardless of whether it is listed in
/etc/fstab or not. Syncing the mount units will not prevent the issue, as it will only synchronize the state of existing mount units with /etc/fstab, not create new ones. Mounting the filesystem manually will not prevent the issue, as it will only mount the filesystem temporarily, not permanently. Remounting all the missing filesystems will not prevent the issue, as it will only mount the filesystems until the next reboot, not after. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Disk Storage, page 457.

NEW QUESTION: 169
After connecting to a remote host via SSH, an administrator attempts to run an application but receives the following error:
[user@workstation ~]$ ssh admin@srv1
Last login: Tue Mar 29 18:03:34 2022
[admin@srvl ~] $ /usr/local/bin/config_manager

Error: cannot open display:
[admin@srv1 ~] $
Which of the following should the administrator do to resolve this error?
A. Disconnect from the SSH session and reconnect using the ssh -x command.
B. Add Options X11 to the /home/admin/.ssh/authorized_keys file.
C. Open port 6000 on the workstation and restart the firewalld service.
D. Enable X11 forwarding in /etc/ssh/ssh_config and restart the server.
Answer: A (LEAVE A REPLY)
Explanation
The error indicates that the application requires an X11 display, but the SSH session does not forward the X11 connection. To enable X11 forwarding, the administrator needs to use the ssh -X option, which requests X11 forwarding with authentication spoofing. This will set the DISPLAY environment variable on the remote host and allow the application to open a window on the local display.
References
CompTIA Linux+ (XK0-005) Certification Study Guide, page 314
Open a window on a remote X display (why "Cannot open display")?, answer by Gilles 'SO- stop being evil'

NEW QUESTION: 170
The security team has identified a web service that is running with elevated privileges A Linux administrator is working to change the systemd service file to meet security compliance standards. Given the following output:

```
[Unit]
Description=CompTIA server daemon
Documentation=man:webserver(8) man:webserver_config(5)
After=network.target

[Service]
Type=notify
EnvironmentFile=/etc/webserver/config
ExecStart=/usr/sbin/webserver -D $OPTIONS
ExecReload=/bin/kill -HUP $MAINPID
KillMode=process
Restart=on-failure
RestartSec=42s

[Install]
WantedBy=multi-user.target
```


Which of the following remediation steps will prevent the web service from running as a privileged user?
A. Removing the ExecStarWusr/sbin/webserver -D SOPTIONS from the service file
B. Updating the Environment File line in the [Service] section to/home/webservice/config
C. Adding the User-webservice to the [Service] section of the service file
D. Changing the:nulti-user.target in the [Install] section to basic.target
Answer: C (LEAVE A REPLY)
Explanation
The remediation step that will prevent the web service from running as a privileged user is adding the User=webservice to the [Service] section of the service file. The service file is a configuration file that defines the properties and behavior of a systemd service. The systemd

is a system and service manager that controls the startup and operation of Linux systems. The service file contains various sections and options that specify how the service should be started, stopped, and managed. The [Service] section defines how the service should be executed and what commands should be run. The User option specifies the user name or ID that the service should run as. The webservice is the name of the user that the administrator wants to run the web service as.
The administrator should add the User=webservice to the [Service] section of the service file, which will prevent the web service from running as a privileged user, such as root, and improve the security of the system. This is the correct remediation step to use to prevent the web service from running as a privileged user. The other options are incorrect because they either do not change the user that the service runs as (removing the ExecStart=/usr/sbin/webserver -D OPTIONS from the service file or updating the EnvironmentFile line in the [Service] section to /home/webservice/config) or do not affect the user that the service runs as (changing the multi-user.target in the [Install] section to basic.target). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing System Services, page 458.

NEW QUESTION: 171
The security team has identified a web service that is running with elevated privileges A Linux administrator is working to change the systemd service file to meet security compliance standards. Given the following output:

Which of the following remediation steps will prevent the web service from running as a privileged user?
A. Removing the ExecStarWusr/sbin/webserver -D SOPTIONS from the service file
B. Updating the Environment File line in the [Service] section to/home/webservice/config
C. Adding the User-webservice to the [Service] section of the service file
D. Changing the:nulti-user.target in the [Install] section to basic.target
Answer: C (LEAVE A REPLY)
Explanation
The remediation step that will prevent the web service from running as a privileged user is adding the User=webservice to the [Service] section of the service file. The service file is a configuration file that defines the properties and behavior of a systemd service. The systemd is a system and service manager that controls the startup and operation of Linux systems. The service file contains various sections and options that specify how the service should be started, stopped, and managed. The [Service] section defines how the service should be executed and what commands should be run. The User option specifies the user name or ID that the service should run as. The webservice is the name of the user that the administrator wants to run the web service as.
The administrator should add the User=webservice to the [Service] section of the service file, which will prevent the web service from running as a privileged user, such as root, and improve the security of the system. This is the correct remediation step to use to prevent

the web service from running as a privileged user. The other options are incorrect because they either do not change the user that the service runs as (removing the ExecStart=/usr/sbin/webserver -D OPTIONS from the service file or updating the EnvironmentFile line in the [Service] section to /home/webservice/config) or do not affect the user that the service runs as (changing the multi-user.target in the [Install] section to basic.target). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing System Services, page 458.

NEW QUESTION: 172
A Linux administrator modified the SSH configuration file. Which of the following commands should be used to apply the configuration changes?
A. systemctl stop sshd
B. systemctl mask sshd
C. systemctl reload sshd
D. systemctl start sshd
Answer: C (LEAVE A REPLY)
The systemctl reload sshd command can be used to apply the configuration changes of the SSH server daemon without restarting it. This is useful to avoid interrupting existing connections. The systemctl stop sshd command would stop the SSH server daemon, not apply the changes. The systemctl mask sshd command would prevent the SSH server daemon from being started, not apply the changes. The systemctl start sshd command would start the SSH server daemon if it is not running, but it would not apply the changes if it is already running. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Secure Shell (SSH), page 415.

NEW QUESTION: 173
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: (SHOW ANSWER)
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: 174
A development team asks an engineer to guarantee the persistency of journal log files across system reboots.
Which of the following commands would accomplish this task?
A. grep -i auto /etc/systemd/journald.conf && systemct1 restart systemd-journald.service
B. cat /etc/systemd/journald.conf | awk '(print $1,$3)'
C. sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i 'persistent/s/'#//q' /etc/systemd/journald.conf
D. journalctl --list-boots && systemct1 restart systemd-journald.service
Answer: C (LEAVE A REPLY)
Explanation
The command sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i 'persistent/s/'#//q' /etc/systemd/journald.conf will accomplish the task of guaranteeing the persistency of journal log files across system reboots. The sed command is a tool for editing text files on Linux systems. The -i option modifies the file in place. The s command substitutes one string for another. The g flag replaces all occurrences of the string. The && operator executes the second command only if the first command succeeds. The q command quits after the first match. The /etc/systemd/journald.conf file is a configuration file for the systemd-journald service, which is responsible for collecting and storing log messages. The command sed -i 's/auto/persistent/g' /etc/systemd/journald.conf will replace the word auto with the word persistent in the file. This will change the value of the Storage option, which controls where the journal log files are stored. The value auto means that the journal log files are stored in the volatile memory and are lost after reboot, while the value persistent means that the journal log files are stored in the persistent storage and are preserved across reboots. The command sed -i 'persistent/s/'#//q' /etc/systemd/journald.conf will remove the # character at the beginning of the line that contains the word persistent. This will uncomment the Storage option and enable it.
The command sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i 'persistent/s/'#//q' /etc/systemd/journald.conf will guarantee the persistency of journal log files across system reboots by changing and enabling the Storage option to persistent. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not change the value of the Storage option (grep -i auto /etc/systemd/journald.conf && systemct1 restart systemd-journald.service or cat /etc/systemd/journald.conf | awk '(print $1,$3)') or do not enable the Storage option (journalctl --list-boots && systemct1 restart systemd-journald.service). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 16: Managing Logging and Monitoring, page 489.

NEW QUESTION: 175
A server is experiencing intermittent connection issues. Some connections to the Internet work as intended, but some fail as if there is no connectivity. The systems administrator inspects the server configuration:

Routing table:

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

IP configuration:

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

ARP table:

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

```
/etc/resolv.conf:
Generated by NetworkHanager
search company.com
nameserver 10.0.5.1
```


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

A. An internal-only DNS server is configured.

B. The IP netmask is wrong for ens3.

C. Two default routes are configured.

D. The ARP table contains incorrect entries.

Answer: C (LEAVE A REPLY)

Explanation

The most likely cause of the issue is that two default routes are configured on the server. The default route is the route that is used when no other route matches the destination of a packet. The default route is usually the gateway that connects the local network to the Internet. The server configuration shows that there are two default routes in the routing table, one with the gateway 192.168.1.1 and the other with the gateway 10.0.0.1.

This can cause a conflict and confusion for the server when deciding which gateway to use for the outgoing packets. Some packets may be sent to the wrong gateway and fail to reach the Internet, while some packets may be sent to the correct gateway and work as intended. This can result in intermittent connection issues and inconsistent behavior. The administrator should remove one of the default routes and keep only the correct one for the network. This can be done by using the ip route del command or by editing the network configuration files. This will resolve the issue and restore the connectivity. The other options are incorrect because they are not supported by the

outputs. The DNS server, the IP netmask, and the ARP table are not the causes of the issue. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, pages 381-382.

NEW QUESTION: 176
Junior system administrator had trouble installing and running an Apache web server on a Linux server. You have been tasked with installing the Apache web server on the Linux server and resolving the issue that prevented the junior administrator from running Apache.
INSTRUCTIONS
Install Apache and start the service. Verify that the Apache service is running with the defaults.
Typing "help" in the terminal will show a list of relevant event commands.
If at any time you would like to bring back the initial state of the simulation, please click the Reset All button.

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

NEW QUESTION: 177
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 is mostly idle, therefore the iowait is high.
B. The system has reached its maximum IOPS, causing the system to be slow.
C. The system has a partitioned disk, which causes the IOPS to be doubled.
D. The system has reached its maximum permitted throughput, therefore iowait is increasing.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 178
A Linux administrator needs to analyze a failing application that is running inside a container. Which of the following commands allows the Linux administrator to enter the running container and analyze the logs that are stored inside?
A. docker run -ti app /bin/sh
B. podman exec -ti app /bin/sh
C. podman run -d app /bin/bash
D. docker exec -d app /bin/bash
Answer: B (LEAVE A REPLY)
Podman exec -ti app /bin/sh allows the Linux administrator to enter the running container and analyze the logs that are stored inside. This command uses the podman tool, which is a daemonless container engine that can run and manage containers on Linux systems. The exec option executes a command inside an existing container, in this case app, which is the name of the container that runs the failing application. The -ti option allocates a pseudo-TTY and keeps STDIN open, allowing for interactive shell access to the container.
The /bin/sh argument specifies the shell command to run inside the container, which can be used to view and manipulate the log files.
The other options are not correct commands for entering a running container and analyzing the logs. Docker run -ti app /bin/sh creates a new container from the app image and runs the /bin/sh command inside it, but does not enter the existing container that runs the failing application. Podman run -d app /bin/bash also creates a new container from the app image and runs the /bin/bash command inside it, but does so in detached mode, meaning that it runs in the background without interactive shell access. Docker exec -d app /bin/bash executes the /bin/bash command inside the existing app container, but also does so in detached mode, without interactive shell access.

NEW QUESTION: 179

The applications team is reporting issues when trying to access the web service hosted in a Linux system. The Linux systems administrator is reviewing the following outputs:
Output 1:
* httpd.service = The Apache HTTPD Server
Loaded: loaded (/usr/lib/systemd/system/httpd.service; disabled; vendor preset: disabled) Active: inactive (dead) Docs: man:httpd(8) man:apachectl(8) Output 2:
16:51:16 up 28 min, 1 user, load average: 0.00, 0.00, 0.07
Which of the following statements best describe the root cause? (Select two).
A. The httpd service is currently started.
B. The httpd service is enabled to auto start at boot time, but it failed to start.
C. The httpd service was manually stopped.
D. The httpd service is not enabled to auto start at boot time.
E. The httpd service runs without problems.
F. The httpd service did not start during the last server reboot.
Answer: C,D (LEAVE A REPLY)
Explanation
The httpd.service is the Apache HTTPD Server, which is a web service that runs on Linux systems. The output 1 shows that the httpd.service is inactive (dead), which means that it is not running. The output 1 also shows that the httpd.service is disabled, which means that it is not enabled to auto start at boot time. Therefore, the statements C and D best describe the root cause of the issue. The statements A, B, E, and F are incorrect because they do not match the output 1. References: [How to Manage Systemd Services on a Linux System]

NEW QUESTION: 180
Which of the following will prevent non-root SSH access to a Linux server?
A. Creating the /etc/nologin file
B. Creating the /etc/nologin.allow file containing only a single line root
C. Creating the /etc/nologin/login.deny file containing a single line +all
D. Ensuring that /etc/pam.d/sshd includes account sufficient pam_nologin.so
Answer: A (LEAVE A REPLY)
Explanation
This file prevents any non-root user from logging in to the system, regardless of the authentication method.
The contents of the file are displayed to the user before the login is terminated. This can be useful for system maintenance or security reasons12.
References: 1: Creating the /etc/nologin File - Oracle 2: How to Restrict Log In Capabilities of Users on Ubuntu

NEW QUESTION: 181
Due to performance issues on a server, a Linux administrator needs to termi-nate an unresponsive process.
Which of the following commands should the
administrator use to terminate the process immediately without waiting for a graceful shutdown?
A. kill -SIGKILL 5545
B. kill -SIGTERM 5545

C. kill -SIGHUP 5545
D. kill -SIGINT 5545
Answer: (SHOW ANSWER)
Explanation
To terminate an unresponsive process immediately without waiting for a graceful shutdown, the administrator can use the command kill -SIGKILL 5545 (A). This will send a signal to the process with the PID 5545 that cannot be ignored or handled by the process, and force it to stop. The other commands will send different signals that may allow the process to perform some cleanup or termination actions, or may be ignored by the process. References:
[CompTIA Linux+ Study Guide], Chapter 6: Managing Processes, Section: Killing Processes
[How to Kill Processes in Linux]

NEW QUESTION: 182
A Linux user reported the following error after trying to connect to the system remotely:
ssh: connect to host 10.0.1.10 port 22: Resource temporarily unavailable The Linux systems administrator executed the following commands in the Linux system while trying to diagnose this issue:

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

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


Which of the following commands will resolve this issue?

A. firewall-cmd --zone=public --permanent --add-service=22

B. systemct1 enable firewalld; systemct1 restart firewalld

C. firewall-cmd --zone=public --permanent --add-service=ssh

D. firewall-cmd --zone=public --permanent --add-port=22/udp

Answer: (SHOW ANSWER)

Explanation

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

NEW QUESTION: 183

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 (LEAVE A REPLY)
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: 184
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: (SHOW ANSWER)
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). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Implementing Basic Security, page 513.

NEW QUESTION: 185

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

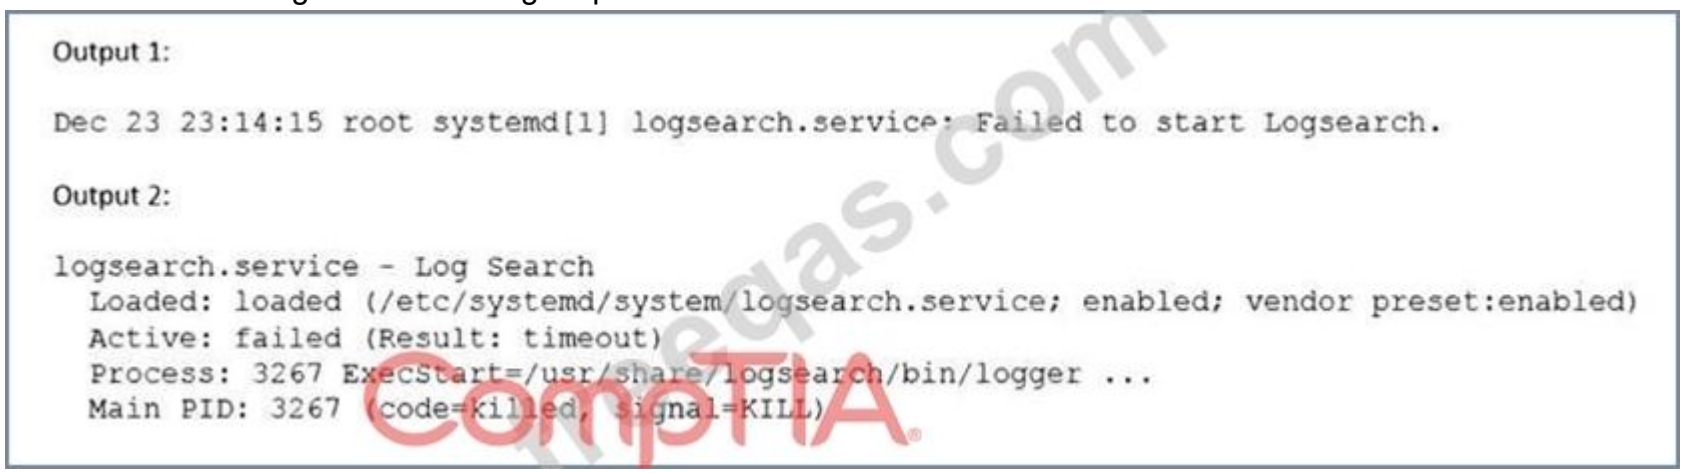

Based on the above outputs, which of the following is the MOST likely action the administrator should take to resolve this issue?

A. Enable the logsearch.service and restart the service.

B. Increase the TimeoutStartUSec configuration for the logsearch.sevice.

C. Update the OnCalendar configuration to schedule the start of the logsearch.service.

D. Update the KillSignal configuration for the logsearch.service to use TERM.

Answer: B (LEAVE A REPLY)

Explanation

The administrator should increase the TimeoutStartUSec configuration for the logsearch.service to resolve the issue. The output of systemct1 status logsearch.service shows that the service failed to start due to a timeout.

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

NEW QUESTION: 186

A systems administrator wants to back up the directory /data and all its contents to /backup/data on a remote server named remote. Which of the following commands will achieve the desired effect?

A. ssh -i /remote:/backup/ /data

B. rsync -a /data remote:/backup/

C. cp -r /data /remote/backup/

D. scp -p /data remote:/backup/data

Answer: (SHOW ANSWER)

NEW QUESTION: 187

A systems administrator has been unable to terminate a process. Which of the following should the administrator use to forcibly stop the process?

A. kill -15

B. kill -1
C. kill -HUP
D. kill -3
E. kill -TERM
Answer: D (LEAVE A REPLY)

NEW QUESTION: 188
A Linux engineer finds multiple failed login entries in the security log file for application users. The Linux engineer performs a security audit and discovers a security issue. Given the following:
grep -iE '*www*|db' /etc/passwd
www-data:x:502:502:www-data:/var/www:/bin/bash
db:x: 505:505:db: /opt/db:/bin/bash
Which of the following commands would resolve the security issue?
A. usermod -d /srv/www-data www-data && usermod -d /var/lib/db db
B. passwd -u www-data && passwd -u db
C. renice -n 1002 -u 502 && renice -n 1005 -u 505
D. chsh -s /bin/false www-data && chsh -s /bin/false db
Answer: D (LEAVE A REPLY)
Explanation
This command will use the chsh tool to change the login shell of the users www-data and db to /bin/false, which means they will not be able to log in to the system1. This will prevent unauthorized access attempts and improve security.
References: 1: Replacing /bin/bash with /bin/false in /etc/passwd file

NEW QUESTION: 189
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
```

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

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

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

D.

Answer: C (LEAVE A REPLY)

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: 190

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

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

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

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


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

A. systemctl unmask mariadb

B. journalctl -g mariadb

C. dnf reinstall mariadb

D. systemctl start mariadb

E. chkconfig mariadb on

F. service mariadb reload

Answer: A,D (LEAVE A REPLY)

Explanation

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

NEW QUESTION: 191
A Linux administrator is adding a new configuration file to a Git repository. Which of the following describes the correct order of Git commands to accomplish the task successfully?
A. pull -> push -> add -> checkout
B. pull -> add -> commit -> push
C. checkout -> push -> add -> pull
D. pull -> add -> push -> commit
Answer: B (LEAVE A REPLY)
Explanation
The correct order of Git commands to add a new configuration file to a Git repository is pull -> add -> commit -> push. The pull command will fetch and merge the changes from the remote repository to the local repository, ensuring that the local repository is up to date. The add command will stage the new configuration file for the next commit, marking it as a new file to be tracked by Git. The commit command will create a new snapshot of the project state with the new configuration file and a descriptive message. The push command will publish the commit to the remote repository, updating the remote branch with the new configuration file.
The pull -> push -> add -> checkout order is incorrect, as it will not create a commit for the new configuration file, and it will switch to a different branch without pushing the changes. The checkout -> push -> add -> pull order is incorrect, as it will switch to a different branch before adding the new configuration file, and it will overwrite the local changes with the remote changes without creating a commit. The pull -> add -> push -> commit order is incorrect, as it will not create a commit before pushing the changes, and it will create a commit that is not synchronized with the remote branch. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 20: Writing and Executing Bash Shell Scripts, page 579.

NEW QUESTION: 192
A Linux user reported the following error after trying to connect to the system remotely:
ssh: connect to host 10.0.1.10 port 22: Resource temporarily unavailable The Linux systems administrator executed the following commands in the Linux system while trying to diagnose this issue:

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

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

Which of the following commands will resolve this issue?

A. firewall-cmd --zone=public --permanent --add-service=22

B. systemct1 enable firewalld; systemct1 restart firewalld

C. firewall-cmd --zone=public --permanent --add-service=ssh

D. firewall-cmd --zone=public --permanent --add-port=22/udp

Answer: (SHOW ANSWER)

Explanation

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

NEW QUESTION: 193

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

Output 1:

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

Output 2:

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


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

A. Update the OnCalendar configuration to schedule the start of the logsearch.service.

B. Update the KillSignal configuration for the logsearch.service to use TERM.

C. Increase the TimeoutStartUSec configuration for the logsearch.sevice.

D. Enable the logsearch.service and restart the service.

Answer: (SHOW ANSWER)

NEW QUESTION: 194

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 (LEAVE A REPLY)

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. Reference: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 13: Networking Fundamentals, page 435.

NEW QUESTION: 195

A Linux administrator needs to resolve a service that has failed to start. The administrator runs the following command:

```
ls -l startup file
```

The following output is returned

```
----------. root root 81k Sep 13 19:01 startupfile
```

Which of the following is MOST likely the issue?

A. The service does not have permissions to read write the startupfile.

B. The service startupfile size cannot be 81k.

C. The service startupfile cannot be owned by root.

D. The service startupfile should not be owned by the root group.

Answer: A (LEAVE A REPLY)

Explanation

The most likely issue is that the service does not have permissions to read or write the startupfile. The output of systemctl status startup.service shows that the service has failed to start and the error message is "Permission denied". The output of ls -l /etc/startupfile shows that the file has the permissions -rw-r--r--, which means that only the owner (root) can read and write the file, while the group (root) and others can only read the file. The service may not run as root and may need write access to the file. The administrator should change the permissions of the file by using the chmod command and grant write access to the group or others, or change the owner or group of the file by using the chown command and assign it to the user or group that runs the service. The other options are incorrect because they are not supported by the outputs. The file size, owner, and group are not the causes of the issue. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, pages 345-346.

NEW QUESTION: 196

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

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


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

A. Execute grub-install --root-directory=/mnt and reboot.

B. Execute grub-install /dev/sdX and reboot.

C. Interrupt the boot process in the GRUB menu and add rescue to the kernel line.

D. Fix the partition modifying /etc/default/grub and reboot.

E. Interrupt the boot process in the GRUB menu and add single to the kernel line.

F. Boot the system on a LiveCD/ISO.

Answer: B,F (LEAVE A REPLY)

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

Boot the system on a LiveCD/ISO. This is necessary to access the system and repair the boot loader. A LiveCD/ISO is a bootable media that contains a Linux distribution that can run without installation. The administrator can boot the system from the LiveCD/ISO and mount the root partition of the system to a temporary directory, such as /mnt.
Execute grub-install /dev/sdX and reboot. This will reinstall the GRUB boot loader to the disk device, where sdX is the device name of the disk, such as sda or sdb. The GRUB boot loader is a program that runs when the system is powered on and allows the user to choose which operating system or kernel to boot. The issue is caused by a corrupted or missing GRUB boot loader, which prevents the system from booting. The command grub-install will restore the GRUB boot loader and fix the issue.
The other options are incorrect because they either do not fix the boot loader (interrupt the boot process in the GRUB menu or fix the partition modifying /etc/default/grub) or do not use the correct syntax (grub-install --root-directory=/mnt instead of grub-install /dev/sdX or rescue or single instead of recovery in the GRUB menu). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing the Linux Boot Process, pages 265-266.

NEW QUESTION: 197
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 -c 2021-09-30 Project_user
B. sudo useradd -m -d 2021-09-30 Project_user
C. sudo modinfo -F 2021-09-30 Project_uses
D. sudo useradd -e 2021-09-30 Project_user
Answer: D (LEAVE A REPLY)

NEW QUESTION: 198
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 (LEAVE A REPLY)
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 .
[CompTIA Linux+ Study Guide], Chapter 11: Working with Containers, Section: Managing Containers with Docker
[Docker PS Command with Examples]

NEW QUESTION: 199
A Linux systems administrator receives a notification that one of the server's filesystems is full. Which of the following commands would help the administrator to identify this filesystem?
A. lsblk
B. fdisk
C. df -h
D. du -ah
Answer: C (LEAVE A REPLY)
Explanation
The df -h command can be used to identify the filesystem that is full. This command displays the disk usage of each mounted filesystem in a human-readable format, showing the total size, used space, available space, and percentage of each filesystem. The lsblk command displays information about block devices, not filesystems.
The fdisk command can be used to manipulate partition tables, not check disk usage. The du -ah command displays the disk usage of each file and directory in a human-readable format, not the filesystems. References: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 14: Managing Disk Storage, page 454.

NEW QUESTION: 200
A Linux administrator has set up a new DNS forwarder and is configuring all internal servers to use the new forwarder to look up external DNS requests. The administrator needs to modify the firewall on the server for the DNS forwarder to allow the internal servers to communicate to it and make the changes persistent between server reboots. Which of the following commands should be run on the DNS forwarder server to accomplish this task?
A. ufw allow out dns
B. systemctl reload firewalld
C. iptables -A OUTPUT -p udp -ra udp -dport 53 -j ACCEPT
D. flrewall-cmd --zone-public --add-port-53/udp --permanent
Answer: (SHOW ANSWER)
Explanation
The command that should be run on the DNS forwarder server to accomplish the task is firewall-cmd --zone=public --add-port=53/udp --permanent. The firewall-cmd command is a tool for managing firewalld, which is a firewall service that provides dynamic and persistent network security on Linux systems. The firewalld uses zones and services to define the rules and policies for the network traffic. The zones are logical groups of network interfaces and sources that have the same level of trust and security. The services are predefined sets of ports and protocols that are associated with certain applications or functions.
The --zone=public option specifies the zone name that the rule applies to. The public zone is the default zone that represents the untrusted network, such as the internet. The --add-port=53/udp option adds a port and protocol to the zone. The 53 is the port number that is used by the DNS service. The udp is the protocol that is used by the DNS service. The --permanent option makes the change persistent across reboots. The command firewall-cmd --zone=public --add-port=53/udp --permanent will modify the firewall on the server for the DNS forwarder to allow the internal servers to communicate to it and make the changes persistent between server reboots. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not modify the firewall on the server for the DNS forwarder (ufw allow out dns or systemctl reload firewalld) or do not use the correct syntax for the command (iptables -A OUTPUT -p udp -ra udp -dport 53 -j ACCEPT instead of iptables -A OUTPUT -p udp -ra udp --dport 53 -j ACCEPT). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 392.

NEW QUESTION: 201
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 (LEAVE A REPLY)
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: 202
Which of the following commands is used to configure the default permissions for new files?
A. chmod
B. sudo
C. setenforce
D. umask
Answer: D (LEAVE A REPLY)

NEW QUESTION: 203
A systems administrator is investigating why one of the servers has stopped connecting to the internet.

```
#curl http://google.com
curl: (6) Could not resolve host: google.com

#cat /etc/resolv.conf
search user.company.com company.com
#nameserver 10.10.10.10

#ip route
0.0.0.0/0 via 10.0.5.1 dev eth0 proto static metric 100
10.0.0.0/16 dev eth0 proto kernel scope link src 10.0.3.60 metric 101

#nmcli connection show
NAME                    UUID                                   TYPE          DEVICE
eth0                    ba4a3d30-efdc-4fa5-83d3-3721fd4aff75   ethernet      eth0
Wired connection 1      8d569d5a-22a2-356d-8532-9a2638f11b5a5   ethernet      --
```


Which of the following is causing the issue?

A. The DNS address has been commented out in the configuration file.

B. The search entry in the /etc/resolv.conf file is incorrect.

C. Wired connection 1 is offline.

D. No default route is defined.

Answer: (SHOW ANSWER)

Explanation

The issue is caused by the lack of a default route defined in the /etc/sysconfig/network-scripts/ifcfg-enp0s3 file. A default route is a special route that specifies where to send packets that do not match any other routes in the routing table. Without a default route, the server will not be able to communicate with hosts outside its local network. The default route is usually configured with the GATEWAY option in the network interface configuration file. For example, to set the default gateway to 192.168.1.1, the file should contain: GATEWAY=192.168.1.1

The other options are not causing the issue. The DNS address is not commented out in the configuration file, it is specified with the DNS1 option. The search entry in the /etc/resolv.conf file is correct, it specifies the domain name to append to unqualified hostnames. Wired connection 1 is online, as indicated by the ONBOOT=yes option and the output of ip link show enp0s3 command. References: Configuring IP Networking with nmcli; Configuring IP Networking with ifcfg Files

NEW QUESTION: 204

A Linux administrator is scheduling a system job that runs a script to check available disk space every hour.

The Linux administrator does not want users to be able to start the job. Given the following:

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

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

[Install]
WantedBy=timers.target
```


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

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

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

A. The checkdiskspace.timer unit should be enabled via systemct1.

B. The timers.target should be reloaded to get the new configuration.

C. The checkdiskspace.timer should be configured to allow manual starts.

D. The checkdiskspace.timer should be started using the sudo command.

Answer: (SHOW ANSWER)

Explanation

The most likely reason the timer will not start is that the checkdiskspace.timer should be configured to allow manual starts. By default, systemd timers do not allow manual activation via systemct1 start, unless they have RefuseManualStart=no in their [Unit] section. This option prevents users from accidentally starting timers that are meant to be controlled by other mechanisms, such as calendar events or dependencies. To enable manual starts for checkdiskspace.timer, the administrator should add RefuseManualStart=no to its [Unit] section and reload systemd.

The other options are not correct reasons for the timer not starting. The checkdiskspace.timer unit does not need to be enabled via systemct1 enable, because enabling a timer only makes it start automatically at boot time or after a system reload, but does not affect manual activation. The timers.target does not need to be reloaded to get the new configuration, because reloading a target only affects units that have a dependency on it, but does not affect manual activation. The checkdiskspace.timer does not need to be started using the sudo command, because the administrator is already running systemct1 as root, as indicated by the # prompt. References: systemd.timer(5) - Linux manual page; systemct1(1) - Linux manual page

NEW QUESTION: 205

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: (SHOW ANSWER)
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: 206
A systems administrator is investigating why one of the servers has stopped connecting to the internet.

```
#curl http://google.com
curl: (6) Could not resolve host: google.com

#cat /etc/resolv.conf
search user.company.com company.com
#nameserver 10.10.10.10

#ip route
0.0.0.0/0 via 10.0.5.1 dev eth0 proto static metric 100
10.0.0.0/16 dev eth0 proto kernel scope link src 10.0.3.60 metric 101

#nmcli connection show
NAME                UUID                                  TYPE          DEVICE
eth0                ba4a3d30-efdc-4fa5-83d3-3721fd4aff75  ethernet      eth0
Wired connection 1  8d569d5a-22a2-356d-8532-9a2638f11b5a5  ethernet      --
```

Which of the following is causing the issue?

A. The DNS address has been commented out in the configuration file.

B. The search entry in the /etc/resolv.conf file is incorrect.

C. Wired connection 1 is offline.

D. No default route is defined.

Answer: D (LEAVE A REPLY)

Explanation

The issue is caused by the lack of a default route defined in the /etc/sysconfig/network-scripts/ifcfg-enp0s3 file. A default route is a special route that specifies where to send packets that do not match any other routes in the routing table. Without a default route, the server will not be able to communicate with hosts outside its local network. The default route is usually configured with the GATEWAY option in the network interface configuration file. For example, to set the default gateway to 192.168.1.1, the file should contain:

GATEWAY=192.168.1.1

The other options are not causing the issue. The DNS address is not commented out in the configuration file, it is specified with the DNS1 option. The search entry in the /etc/resolv.conf file is correct, it specifies the domain name to append to unqualified hostnames. Wired connection 1 is online, as indicated by the ONBOOT=yes option and the output of ip link show enp0s3 command. References: Configuring IP Networking with nmcli; Configuring IP Networking with ifcfg Files

NEW QUESTION: 207

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

A. podman exec -ti app /bin/sh

B. docker exec -d app /bin/bash

C. docker run -ti app /bin/sh

D. podman run -d app /bin/bash

Answer: A (LEAVE A REPLY)

NEW QUESTION: 208

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

$ sudo visudo

>>> /etc/sudoers: syntax error near line 28 <<<

sudo: parse error in /etc/sudoers near line 28

sudo: no valid sudoers sources found, quitting

The following output is provided:

grep root /etc/shadow

root :* LOCK *: 14600 ::::::

Which of the following actions will resolve this issue?

A. Comment out line 28 from /etc/sudoers and try to use sudo again.

B. Log in to the system using the other regular user, switch to root, and comment out line 28 from /etc/sudoers.

C. Boot the system in single user mode and comment out line 28 from /etc/sudoers.

D. Log in directly using the root account and comment out line 28 from /etc/sudoers.

Answer: (SHOW ANSWER)

NEW QUESTION: 209

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

```
$ vmstat -s --unit M

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

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

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

Which of the following commands would address the issue?

A. kill -9 8321

B. renice -10 8321

C. free 8321

D. top -p 8321

Answer: B (LEAVE A REPLY)

NEW QUESTION: 210

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: (SHOW ANSWER)

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: 211
A Linux administrator is troubleshooting a memory-related issue. Based on the output of the commands:

```
$ vmstat -s --unit M

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

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

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


Which of the following commands would address the issue?
A. top -p 8321
B. kill -9 8321
C. renice -10 8321
D. free 8321
Answer: B (LEAVE A REPLY)
Explanation

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

NEW QUESTION: 212
A new disk was presented to a server as /dev/ sdd. The systems administrator needs to check if a partition table is on that disk. Which of the following commands can show this information?
A. lsscsi
B. fdisk
C. blkid
D. partprobe
Answer: B (LEAVE A REPLY)
Explanation
The command that can be used to check if a partition table is on a disk is fdisk. The fdisk command can display, create, delete, and modify partitions on a disk. To show the partition table of a disk, the administrator can use fdisk -l /dev/sdd (B). References:
[CompTIA Linux+ Study Guide], Chapter 5: Managing Filesystems and Logical Volumes, Section:
Partitioning Disks
[How to Use Fdisk Command in Linux]

NEW QUESTION: 213
A Linux administrator is trying to remove the ACL from the file /home/user/data. txt but receives the following error message:

```
setfacl: data.txt: operation not permitted
```

Given the following analysis:

```
/dev/mapper/linux-home on /home type xfs (rw,relatime,seclabel,attr2,inode64,usrquota)

-rw-rw-r--+ 1 user staff 2354 Sep 15 16:33 data.txt
-rw-rw-r--+ user staff unconfined_u:object_r:user_home_t:s0 data.txt

# file: data.txt
# owner: user
# group: staff
user::rw-
user:accounting:rw-
group::r-
mask::rw-
other::r-

Attributes:
-----a---------
```


Which of the following is causing the error message?

A. The administrator is not using a highly privileged account.

B. The filesystem is mounted with the wrong options.

C. SELinux file context is denying the ACL changes.

D. File attributes are preventing file modification.

Answer: D (LEAVE A REPLY)

Explanation

File attributes are preventing file modification, which is causing the error message. The output of lsattr /home/user/data.txt shows that the file has the immutable attribute (i) set, which means that the file cannot be changed, deleted, or renamed. The command setfacl -b /home/user/data.txt tries to remove the ACL from the file, but fails because of the immutable attribute. The administrator needs to remove the immutable attribute first by using the command chattr -i /home/user/data.txt and then try to remove the ACL again. The other options are incorrect because they are not supported by the outputs. The administrator is using a highly privileged account, as shown by the # prompt. The filesystem is mounted with the correct options, as shown by the output of mount | grep /home. SELinux file context is not denying the ACL changes, as shown by the output of ls -Z /home/user/data.txt. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 11: Managing Files and Directories, pages 357-358.

NEW QUESTION: 214

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

```
Device mismatch detected
```

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

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

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

A. mount disk by-label

B. mount disk by device-id

C. mount disk by-blkid

D. fsck -A

Answer: C (LEAVE A REPLY)

NEW QUESTION: 215

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

A. pull -> push -> add -> checkout

B. checkout -> push -> add -> pull

C. pull -> add -> push -> commit

D. pull -> add -> commit -> push

Answer: D (LEAVE A REPLY)

NEW QUESTION: 216

A Linux administrator needs to resolve a service that has failed to start. The administrator runs the following command:

```
ls -l startup file
```

The following output is returned

```
----------. root root 81k Sep 13 19:01 startupfile
```

Which of the following is MOST likely the issue?

A. The service startupfile size cannot be 81k.

B. The service startupfile cannot be owned by root.

C. The service startupfile should not be owned by the root group.

D. The service does not have permissions to read write the startupfile.

Answer: (SHOW ANSWER)

NEW QUESTION: 217

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_rsa
B. id_ecdsa
C. id_rsa.pub
D. id_dsa.pem
Answer: (SHOW ANSWER)

NEW QUESTION: 218
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.sh
chmod 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 (LEAVE A REPLY)
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: 219
What is the main objective when using Application Control?
A. To assist the firewall blade with handling traffic.
B. To filter out specific content.
C. Ensure security and privacy of information.
D. To see what users are doing.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 220
A Linux administrator needs to create an image named sda.img from the sda disk and store it in the /tmp directory. Which of the following commands should be used to accomplish this task?
A. dd of=/dev/sda if=/tmp/sda.img
B. dd if=/dev/sda of=/tmp/sda.img
C. dd --of=/dev/sda --if=/tmp/sda.img
D. dd --if=/dev/sda --of=/tmp/sda.img
Answer: B (LEAVE A REPLY)

NEW QUESTION: 221
A Linux administrator is scheduling a system job that runs a script to check available disk space every hour.
The Linux administrator does not want users to be able to start the job. Given the following:

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

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

[Install]
WantedBy=timers.target
```


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

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

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

NEW QUESTION: 222
A systems administrator is tasked with mounting a USB drive on a system. The USB drive has a single partition, and it has been mapped by the system to the device /dev/sdb. Which of the following commands will mount the USB to /media/usb?
A. mount /dev/sdb1 /media/usb
B. mount /dev/sdb0 /media/usb
C. mount /dev/sdb /media/usb
D. mount -t usb /dev/sdb1 /media/usb
Answer: (SHOW ANSWER)
The mount /dev/sdb1 /media/usb command will mount the USB drive to /media/usb. This command will attach the filesystem on the first partition of the USB drive (/dev/sdb1) to the mount point /media/usb, making it accessible to the system. The mount /dev/sdb0 /media/usb command is invalid, as there is no such device as /dev/sdb0. The mount /dev/sdb /media/usb command is incorrect, as it will try to mount the entire USB drive instead of its partition, which may cause errors or data loss. The mount -t usb /dev/sdb1 /media/usb command is incorrect, as usb is not a valid filesystem type for mount. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 14: Managing Disk Storage, page 455.

NEW QUESTION: 223
A systems administrator configured firewall rules using firewalld. However, after the system is rebooted, the firewall rules are not present:

```
Chain INPUT (policy ACCEPT)
target     prot opt source               destination

Chain FORWARD (policy ACCEPT)
target     prot opt source               destination

Chain OUTPUT (policy ACCEPT)
target     prot opt source               destination
```

The systems administrator makes additional checks:

```
- dynamic firewall daemon
  Loaded: loaded (/usr/lib/systemd/system/firewalld.service; disabled; vendor preset: enabled)
  Active: inactive (dead)
  Docs: man: firewalld (1)

firewalld is not running
```

Which of the following is the reason the firewall rules are not active?

A. iptables is conflicting with firewalld.

B. The wrong system target is activated.

C. FIREWALL_ARGS has no value assigned.

D. The firewalld service is not enabled.

Answer: D (LEAVE A REPLY)

Explanation

The reason the firewall rules are not active is that the firewalld service is not enabled. This means that the service will not start automatically at boot time or after a system reload. To enable the firewalld service, the systems administrator needs to use the command sudo systemct1 enable firewalld. This will create a symbolic link from the firewalld service file to the appropriate systemd target, such as multi-user.target. Enabling the service does not start it immediately, so the systems administrator also needs to use the command sudo systemct1 start firewalld or sudo systemct1 reload firewalld to activate the firewall rules.

The other options are not correct reasons for the firewall rules not being active. iptables is not conflicting with firewalld, because firewalld uses iptables as its backend by default. The wrong system target is not activated, because firewalld is independent of the system target and can be enabled for any target. FIREWALL_ARGS has no value assigned, but this is not a problem, because FIREWALL_ARGS is an optional environment variable that can be used to pass additional arguments to the firewalld daemon, such as --debug or --nofork. If FIREWALL_ARGS is empty or not defined, firewalld will use its default arguments. References: firewalld.service(8) - Linux manual page; firewall-cmd(1) - Linux manual page; systemct1(1) - Linux manual page

NEW QUESTION: 224

A systems engineer has deployed a new application server, but the server cannot communicate with the backend database hostname. The engineer confirms that the application server can ping the database server's IP address. Which of the following is the most likely cause of the issue?

A. Incorrect DNS servers

B. Unreachable default gateway

C. Missing route configuration

D. Misconfigured subnet mask

Answer: (SHOW ANSWER)

Explanation

This is because the application server can ping the database server's IP address, but not its hostname, which suggests that the DNS resolution is not working properly. DNS servers are responsible for translating hostnames into IP addresses, and vice versa. If the application server has incorrect or unreachable DNS servers configured, it will not be able to resolve the hostname of the database server and communicate with it.

To troubleshoot this issue, the systems engineer should check the DNS configuration on the application server, which is usually stored in the /etc/resolv.conf file. This file should contain valid nameserver entries that point to the DNS servers that can resolve the database server's hostname. 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.

Alternatively, the systems engineer can use the nslookup or dig commands to test the DNS resolution of the database server's hostname from the application server. These commands will query a specified DNS server and return the IP address of the hostname, or an error message if the resolution fails. For example, to query Google's public DNS server for the IP address of comptia.org, the command would be:

nslookup comptia.org 8.8.8.8

or

dig comptia.org @8.8.8.8

NEW QUESTION: 225

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: (SHOW ANSWER)

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

NEW QUESTION: 226

A systems administrator pressed Ctrl+Z after starting a program using the command line, and the shell prompt was presented. In order to go back to the program, which of the following commands can the administrator use?

A. fg

B. su

C. bg
D. ed
Answer: (SHOW ANSWER)
Ctrl+Z suspended the process, and "fg" will bring it back into the foreground of the shell A Comprehensive and Detailed To go back to a program that was suspended by pressing Ctrl+Z in the command line, the command that can be used is fg. The fg command stands for foreground, and it resumes the job that is next in the queue and brings it to the foreground. Alternatively, if there are more than one suspended jobs, fg can be followed by a job number to resume a specific job. The other commands are incorrect because they either do not resume a suspended job, or they have different functions such as switching user (su), pushing a job to the background (bg), or editing a file (ed). Reference: CompTIA Linux+ Study Guide, Fourth Edition, page 181-182.

NEW QUESTION: 227
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: (SHOW ANSWER)
Explanation
The command docker rm $(docker ps -aq) will allow the administrator to clean up the containers in an exited state. The docker command is a tool for managing Docker containers on Linux systems. Docker containers are isolated and lightweight environments that can run applications and services without affecting the host system.
Docker uses images to create containers, which are files that contain the code, libraries, dependencies, and configuration of the applications and services. The rm option removes one or more containers. The $(docker ps -aq) is a command substitution that executes the command inside the parentheses and replaces it with the output. The docker ps -aq command lists all the containers, including the ones in an exited state, and shows only their IDs. The docker rm $(docker ps -aq) command will remove all the containers, including the ones in an exited state, by passing their IDs to the rm option. This will allow the administrator to clean up the containers in an exited state. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not exist (docker rm --all or docker rm --state exited) or do not remove the containers (docker images prune *). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Cloud and Virtualization Technologies, page 571.

NEW QUESTION: 228
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: (SHOW ANSWER)
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: 229
A systems administrator is investigating a service that is not starting up. Given the following information:

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

Which of the following systemd commands should the administrator use in order to obtain more details about the failing service?
A. systemct1 analyze network
B. systemct1 info network
C. sysctl -a network
D. journalctl -xu network
Answer: D (LEAVE A REPLY)
Explanation
The systemd is a system and service manager for Linux systems that provides a standard way to control and monitor system services. The systemd uses various commands and tools to manage and troubleshoot system services, such as systemct1, sysctl, and journalctl. The systemct1 command is used to start, stop, enable, disable, restart, reload, status, and list system services. The sysctl command is used to configure kernel parameters at runtime. The journalctl command is used to view and filter the logs of system services.
To investigate a service that is not starting up, the administrator can use the journalctl command with the -xu option. The -x option enables verbose output that includes explanatory text and priority information.
The -u option filters the output by a specific unit name, such as network.service. Therefore, the command journalctl -xu network will show detailed logs of the network service, which can help identify the cause of the failure. The statement D is correct.
The statements A, B, and C are incorrect because they do not provide more details about the failing service.

The systemct1 analyze network command does not exist. The systemct1 info network command shows basic information about the network unit, such as description, load state, active state, sub state, and main PID.
The sysctl -a network command shows all kernel parameters related to network settings. References: [How to Use Systemd to Manage System Services]

NEW QUESTION: 230
A Linux administrator was asked to run a container with the httpd server inside. This container should be exposed at port 443 of a Linux host machine while it internally listens on port 8443. Which of the following commands will accomplish this task?
A. podman run -d -p 8443:443 httpd
B. podman run -d -p 443:8443 httpd
C. podman exec -p 8443:443 httpd
D. podman run -d -e 443:8443 httpd
Answer: (SHOW ANSWER)

NEW QUESTION: 231
A Linux user reported the following error after trying to connect to the system remotely:
ssh: connect to host 10.0.1.10 port 22: Resource temporarily unavailable The Linux systems administrator executed the following commands in the Linux system while trying to diagnose this issue:

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

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


Which of the following commands will resolve this issue?
A. firewall-cmd --zone=public --permanent --add-service=22
B. systemct1 enable firewalld; systemct1 restart firewalld
C. firewall-cmd --zone=public --permanent --add-service=ssh

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

NEW QUESTION: 232
A junior Linux administrator is tasked with installing an application. The installation guide states the application should only be installed in a run level 5 environment.

```
$ systemctl get-default
getty.target
```

Which of the following commands would ensure the server is set to runlevel 5?
A. systemctl isolate multi-user.target
B. systemctl isolate graphical.target
C. systemctl isolate network.target
D. systemctl isolate basic.target
Answer: B (LEAVE A REPLY)
Explanation
The command that would ensure the server is set to runlevel 5 is systemctl isolate graphical.target. This command will change the current target (or runlevel) of systemd to graphical.target, which is equivalent to runlevel 5 in SysV init systems. Graphical.target means that the system will start with a graphical user interface (GUI) and all services required for it.
The other options are not correct commands for setting the server to runlevel 5. The systemctl isolate multi-user.target command will change the current target to multi-user.target, which is equivalent to runlevel 3 in SysV init systems. Multi-user.target means that the system will start with multiple user logins and networking, but without a GUI. The systemctl isolate network.target command will change the current target to network.target, which is not a real runlevel but a synchronization point for network-related services.
Network.target means that network functionality should be available, but does not specify whether it should be started before or after it.
The systemctl isolate basic.target command will change the current target to basic.target, which is also not a real runlevel but a synchronization point for basic system services. Basic.target means that all essential services should be started, but does not specify whether it should be started before or after it. References: systemd System and Service Manager; systemd.special(7) - Linux manual page

NEW QUESTION: 233
A systems administrator created a web server for the company and is required to add a tag for the API so end users can connect. Which of the following would the administrator do to complete this requirement?
A. hostnamectl status --no-ask-password

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

NEW QUESTION: 234
A Linux administrator is creating a primary partition on the replacement hard drive for an application server. Which of the following commands should the administrator issue to verify the device name of this partition?
A. sudo fdisk /dev/sda
B. sudo fdisk -s /dev/sda
C. sudo fdisk -l
D. sudo fdisk -h
Answer: C (LEAVE A REPLY)
The command sudo fdisk -l should be issued to verify the device name of the partition. The sudo command allows the administrator to run commands as the superuser or another user. The fdisk command is a tool for manipulating disk partitions on Linux systems. The -l option lists the partitions on all disks or a specific disk. The command sudo fdisk -l will show the device names, sizes, types, and other information of the partitions on all disks. The administrator can identify the device name of the partition by looking at the output. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not list the partitions (sudo fdisk /dev/sda or sudo fdisk -h) or do not exist (sudo fdisk -s /dev/sda). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 317.

NEW QUESTION: 235
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 (LEAVE A REPLY)
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: 236
An administrator attempts to rename a file on a server but receives the following error.

```
mv: cannot move 'files/readme.txt' to 'files/readme.txt.orig': Operation not permitted.
```

The administrator then runs a few commands and obtains the following output:

```
$ ls -ld files/
drwxrwxrwt.1    users    users    20    Sep 10 15:15    files/

$ ls -a files/
drwxrwxrwt.1    users    users    20    Sep 10 15:15    -
drwxr-xr-x.1    users    users    32    Sep 10 15:15    ..
-rw-rw-r--.1    users    users    4     Sep 12 10:34    readme.txt
```


Which of the following commands should the administrator run NEXT to allow the file to be renamed by any user?
A. chgrp reet files
B. chacl -R 644 files
C. chown users files
D. chmod -t files

Answer: D (LEAVE A REPLY)

Explanation

The command that the administrator should run NEXT to allow the file to be renamed by any user is chmod -t files. This command uses the chmod tool, which is used to change file permissions and access modes. The -t option removes (or sets) the sticky bit on a directory, which restricts deletion or renaming of files within that directory to only their owners or root. In this case, since files is a directory with sticky bit set (indicated by t in drwxrwxrwt), removing it will allow any user to rename or delete files within that directory.

The other options are not correct commands for allowing any user to rename files within files directory. The chgrp reet files command will change the group ownership of files directory to reet, but it will not affect its permissions or access modes. The chacl -R 644 files command is invalid, as chacl is used to change file access control lists (ACLs), not permissions or access modes. The chown users files command will change the user ownership of files directory to users, but it will not affect its permissions or access modes. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing Users and Groups; chmod(1) - Linux manual page

NEW QUESTION: 237

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

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


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

A. chgrp -R 755 data/

B. chmod -R 777 data/

C. chattr -R -i data/

D. chown -R data/

Answer: C (LEAVE A REPLY)

The command that can be used to resolve the issue of being unable to remove a particular data folder is chattr -R -i data/. This command will use the chattr utility to change file attributes on a Linux file system. The -R option means that chattr will recursively change attributes of directories and their contents. The -i option means that chattr will remove (unset) the immutable attribute from files or directories. When a file or directory has the immutable attribute set, it cannot be modified, deleted, or renamed.

The other options are not correct commands for resolving this issue. The chgrp -R 755 data/ command will change the group ownership of data/ and its contents recursively to 755, which is not a valid group name. The chgrp command is used to change group ownership of files or directories. The chmod -R 777 data/ command will change the file mode bits of data/ and its contents recursively to 777, which means

that everyone can read, write, and execute them. However, this will not remove the immutable attribute, which prevents deletion or modification regardless of permissions. The chmod command is used to change file mode bits of files or directories. The chown -R data/ command is incomplete and will produce an error. The chown command is used to change the user and/or group ownership of files or directories, but it requires at least one argument besides the file name. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 7: Managing Disk Storage; chattr(1) - Linux manual page; chgrp(1) - Linux manual page; chmod(1) - Linux manual page; chown(1) - Linux manual page

NEW QUESTION: 238
After listing the properties of a system account, a systems administrator wants to remove the expiration date of a user account. Which of the following commands will accomplish this task?
A. chgrp system accountname
B. passwd -s accountname
C. chmod -G system account name
D. chage -E -1 accountname
Answer: (SHOW ANSWER)
Explanation
The command chage -E -1 accountname will accomplish the task of removing the expiration date of a user account. The chage command is a tool for changing user password aging information on Linux systems.
The -E option sets the expiration date of the user account, and the -1 value means that the account will never expire. The command chage -E -1 accountname will remove the expiration date of the user account named accountname. This is the correct command to use to accomplish the task. The other options are incorrect because they either do not affect the expiration date (chgrp, passwd, or chmod) or do not exist (chmod
-G). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15: Managing Users and Groups, page 467.

NEW QUESTION: 239
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: A,E (LEAVE A REPLY)
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: 240
A systems administrator received a notification that a system is performing slowly. When running the top command, the systems administrator can see the following values:

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

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

NEW QUESTION: 241
A Linux administrator needs to create an image named sda.img from the sda disk and store it in the /tmp directory. Which of the following commands should be used to accomplish this task?
A. dd of=/dev/sda if=/tmp/sda.img
B. dd if=/dev/sda of=/tmp/sda.img
C. dd --if=/dev/sda --of=/tmp/sda.img
D. dd --of=/dev/sda --if=/tmp/sda.img
Answer: B (LEAVE A REPLY)

The command dd if=/dev/sda of=/tmp/sda.img should be used to create an image named sda.img from the sda disk and store it in the /tmp directory. The dd command is a tool for copying and converting data on Linux systems. The if option specifies the input file or device, in this case /dev/sda, which is the disk device. The of option specifies the output file or device, in this case /tmp/sda.img, which is the image file. The command dd if=/dev/sda of=/tmp/sda.img will copy the entire disk data from /dev/sda to /tmp/sda.img and create an image file. This is the correct command to use to accomplish the task. The other options are incorrect because they either use the wrong options (--if or --of instead of if or of) or swap the input and output (dd of=/dev/sda if=/tmp/sda.img or dd --of=/dev/sda --if=/tmp/sda.img). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, page 323.

NEW QUESTION: 242
A Linux administrator cloned an existing Linux server and built a new server from that clone. The administrator encountered the following error after booting the cloned server:

```
Device mismatch detected
```

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

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

Which of the following should administrator use to resolve the device mismatch issue and mount the disk?
A. mount disk by device-id
B. fsck -A
C. mount disk by-label
D. mount disk by-blkid
Answer: A (LEAVE A REPLY)
The administrator should use the command mount disk by device-id to resolve the device mismatch issue and mount the disk. The issue is caused by the cloned server having a different device name for the disk than the original server. The output of blkid shows that the disk has the device name /dev/sdb1 on the cloned server, but the output of cat /etc/fstab shows that the disk is expected to have the device name /dev/sda1. The command mount disk by device-id will mount the disk by using its unique identifier (UUID) instead of its device name. The UUID can be obtained from the output of blkid or lsblk -f. The command will mount the disk to the specified mount point (/data) and resolve the issue. The other options are incorrect because they either do not mount the disk (fsck -A), do not use the correct identifier

(mount disk by-label or mount disk by-blkid), or do not exist (mount disk by-blkid). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 318-319.

NEW QUESTION: 243
A developer is trying to install an application remotely that requires a graphical interface for installation. The developer requested assistance to set up the necessary environment variables along with X11 forwarding in SSH. Which of the following environment variables must be set in remote shell in order to launch the graphical interface?
A. $RHOST
B. SETENV
C. $SHELL
D. $DISPLAY
Answer: (SHOW ANSWER)
The environment variable that must be set in remote shell in order to launch the graphical interface is $DISPLAY. This variable tells X11 applications where to display their windows on screen. It usually has the form hostname:displaynumber.screennumber, where hostname is the name of the computer running the X server, displaynumber is a unique identifier for an X display on that computer, and screennumber is an optional identifier for a screen within an X display. For example, localhost:0.0 means display number 0 on the local host. If the hostname is omitted, it defaults to the local host.
The other options are not correct environment variables for launching the graphical interface. $RHOST is a variable that stores the name of the remote host, but it is not used by X11 applications. SETENV is a command that sets environment variables in some shells, but it is not an environment variable itself. $SHELL is a variable that stores the name of the current shell, but it is not related to X11 forwarding.
Reference: How to enable or disable X11 forwarding in an SSH server; How to Configure X11 Forwarding Using SSH In Linux

NEW QUESTION: 244
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 (LEAVE A REPLY)
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: 245
An administrator created an initial Git repository and uploaded the first files. The administrator sees the following when listing the repository:

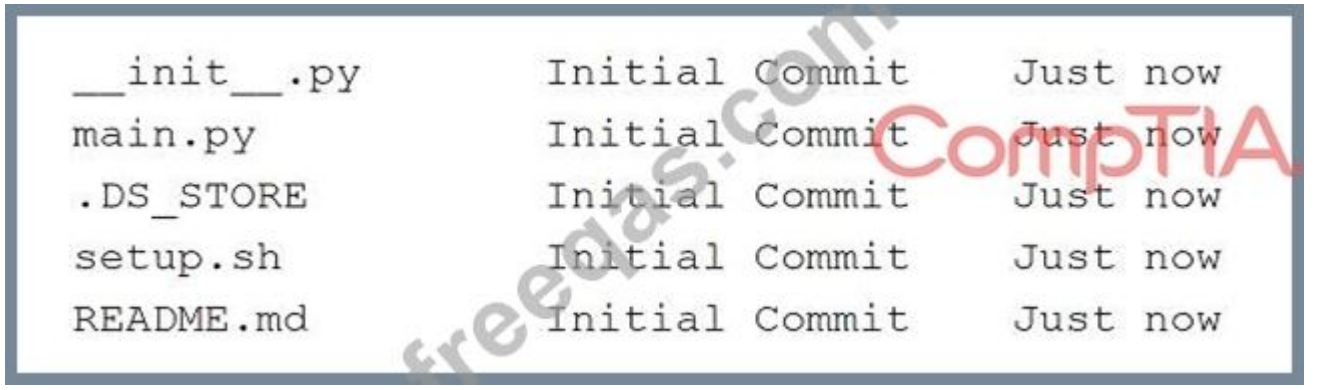

The administrator notices the file . DS STORE should not be included and deletes it from the online repository. Which of the following should the administrator run from the root of the local repository before the next commit to ensure the file is not uploaded again in future commits?

A. rm -f .DS STORE && git push

B. git fetch && git checkout .DS STORE

C. rm -f .DS STORE && git rebase origin main

D. echo .DS STORE >> .gitignore

Answer: D (LEAVE A REPLY)

Explanation

The correct answer is D. The administrator should run "echo .DS STORE >> .gitignore" from the root of the local repository before the next commit to ensure the file is not uploaded again in future commits.

This command will append the file name .DS STORE to the end of the .gitignore file, which is a special file that tells Git to ignore certain files or directories that should not be tracked or uploaded to the repository. By adding .DS STORE to the .gitignore file, the administrator will prevent Git from staging, committing, or pushing this file in the future.

The other options are incorrect because:

A: rm -f .DS STORE && git push

This command will delete the file .DS STORE from the local repository and then push the changes to the remote repository. However, this does not prevent the file from being uploaded again in future commits, if it is recreated or copied to the local repository.

B: git fetch && git checkout .DS STORE

This command will fetch the latest changes from the remote repository and then restore the file .DS STORE from the remote repository to the local repository. This is not what the administrator wants to do, as this will undo the deletion of the file from the online repository.

C: rm -f .DS STORE && git rebase origin main

This command will delete the file .DS STORE from the local repository and then rebase the local branch onto the main branch of the remote repository. This will rewrite the commit history of the local branch and may cause conflicts or errors. This is not what the administrator wants to do, as this is a risky and unnecessary operation.

NEW QUESTION: 246

The development team created a new branch with code changes that a Linux administrator needs to pull from the remote repository. When the administrator looks for the branch in Git, the branch in question is not visible.

Which of the following commands should the Linux administrator run to refresh the branch information?

A. git fetch

B. git checkout

C. git clone

D. git branch

Answer: A (LEAVE A REPLY)

Explanation

The git fetch command downloads commits, files, and refs from a remote repository into the local one. It also updates the remote-tracking branches, which are references to the state of the remote branches. By running git fetch, the administrator can see the new branch created by the development team and then use git checkout to switch to it12.

References: 1: Git - git-fetch Documentation 2: Git Fetch | Atlassian Git Tutorial

NEW QUESTION: 247

When trying to log in remotely to a server, a user receives the following message:

```
Password:
Last failed login: Wed Sep 15 17:23:45 CEST 2021 from 10.0.4.3 on ssh:notty
There were 3 failed login attempts since the last successful login.
Connection to localhost closed.
```

The server administrator is investigating the issue on the server and receives the following outputs:

Output 1:

```
user:x:1001:1374::/home/user:/bin/false
```

Output 2:

```
dzwx------. 2 user 62 Sep 15 17:17 /home/user
```

Output 3:

```
Sep 12 14:14:05 server sshd[22958]: Failed password for user from 10.0.2.8
Sep 15 17:24:03 server sshd[8460]: Accepted keyboard-interactive/pam for user from 10.0.6.5 port 50928 ssh2
Sep 15 17:24:03 server sshd[8460]: pam_unix(sshd:session): session opened for user testuser
Sep 15 17:24:03 server sshd[8460]: pam_unix(sshd:session): session closed for user testuser
```

Which of the following is causing the issue?

A. The wrong permissions are on the user's home directory.

B. The account was locked out due to three failed logins.

C. The user entered the wrong password.

D. The user has the wrong shell assigned to the account.

Answer: D (LEAVE A REPLY)

The user has the wrong shell assigned to the account, which is causing the issue. The output 1 shows that the user's shell is set to /bin/false, which is not a valid shell and will prevent the user from logging in. The output 2 shows that the user's home directory has the correct permissions (drwxr-xr-x), and the output 3 shows that the user entered the correct password and was accepted by the SSH daemon, but the session was closed immediately due to the invalid shell. The other options are incorrect because they are not supported by the outputs. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 413.

NEW QUESTION: 248

A new file was added to a main Git repository. An administrator wants to synchronize a local copy with the contents of the main repository. Which of the following commands should the administrator use for this task?

A. git pull

B. git push

C. git reflog

D. git status

Answer: A (LEAVE A REPLY)

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: 249

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/script. sh

Answer: A (LEAVE A REPLY)

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: 250

A new application container was built with an incorrect version number. Which of the following commands should be used to rename the image to match the correct version 2.1.2?

A. docker tag comptia/app:2.1.1 comptia/app:2.1.2

B. docker push comptia/app:2.1.1 comptia/app:2.1.2

C. docker rmi comptia/app:2.1.1 comptia/app:2.1.2

D. docker update comptia/app:2.1.1 comptia/app:2.1.2

Answer: A (LEAVE A REPLY)
Explanation
The best command to use to rename the image to match the correct version 2.1.2 is A. docker tag comptia/app:2.1.1 comptia/app:2.1.2. This command will create a new tag for the existing image with the new version number, without changing the image content or ID. The other commands are either incorrect or not suitable for this task. For example:
B: docker push comptia/app:2.1.1 comptia/app:2.1.2 will try to push two images to a remote repository, but it does not rename the image locally.
C: docker rmi comptia/app:2.1.1 comptia/app:2.1.2 will try to remove two images from the local system, but it does not rename the image.
D: docker update comptia/app:2.1.1 comptia/app:2.1.2 will try to update the configuration of a running container, but it does not rename the image.

NEW QUESTION: 251
Which of the following tools is commonly used for creating CI/CD pipelines?
A. Chef
B. Puppet
C. Jenkins
D. Ansible
Answer: C (LEAVE A REPLY)
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: 252
After installing a new version of a package, a systems administrator notices a new version of the corresponding, service file was Installed In order to use the new version of the, service file, which of the following commands must be Issued FIRST?
A. systemctl status
B. systemctl stop
C. systemctl reinstall
D. systemctl daemon-reload
Answer: D (LEAVE A REPLY)
Explanation
After installing a new version of a package that includes a new version of the corresponding service file, the systemctl daemon-reload command must be issued first in order to use the new version of the service file. This command will reload the systemd manager configuration and read all unit files that have changed on disk. This will ensure that systemd recognizes the new service file and applies its settings correctly. The systemctl status command will display information about a service unit, but it will not reload the configuration. The systemctl stop command will stop a service unit, but it will not reload the configuration. The systemctl reinstall command does not exist.
References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17:
System Maintenance and Operation, page 518.

NEW QUESTION: 253
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 (LEAVE A REPLY)
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: 254
A Linux administrator was asked to run a container with the httpd server inside. This container should be exposed at port 443 of a Linux host machine while it internally listens on port 8443. Which of the following commands will accomplish this task?
A. podman run -d -p 443:8443 httpd
B. podman run -d -p 8443:443 httpd
C. podman run -d -e 443:8443 httpd
D. podman exec -p 8443:443 httpd
Answer: A (LEAVE A REPLY)
The command that will accomplish the task of running a container with the httpd server inside and exposing it at port 443 of the Linux host machine while it internally listens on port 8443 is podman run -d -p 443:8443 httpd. This command uses the podman tool, which is a daemonless container engine that can run and manage containers on Linux systems. The -d option runs the container in detached mode, meaning that it runs in the background without blocking the terminal. The -p option maps a port on the host machine to a port inside the container, using the format host_port:container_port. In this case, port 443 on the host machine is mapped to port 8443 inside the container, allowing external access to the httpd server. The httpd argument specifies the name of the image to run as a container, which in this case is an image that contains the Apache HTTP Server software. The other options are not correct commands for accomplishing the task. Podman run -d -p 8443:443 httpd maps port 8443 on the host machine to port 443 inside the container, which does not match the requirement. Podman run -d -e 443:8443 httpd uses the -e option instead of the -p option, which sets an environment variable inside the container instead of mapping a port. Podman exec -p 8443:443 httpd uses the podman exec command instead of the podman run command, which executes a command inside an existing container instead of creating a new one. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 18: Automating Tasks

NEW QUESTION: 255
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: (SHOW ANSWER)
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: 256
A junior administrator is trying to set up a passwordless SSH connection to one of the servers. The administrator follows the instructions and puts the key in the authorized_key file at the server, but the administrator is still asked to provide a password during the connection. Given the following output:

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

Which of the following commands would resolve the issue and allow an SSH connection to be established without a password?
A. restorecon -rv .ssh/authorized_key
B. mv .ssh/authorized_key .ssh/authorized_keys
C. systemct1 restart sshd.service
D. chmod 600 mv .ssh/authorized_key
Answer: B (LEAVE A REPLY)
Explanation
The command mv .ssh/authorized_key .ssh/authorized_keys will resolve the issue and allow an SSH connection to be established without a password. The issue is caused by the incorrect file name of the authorized key file on the server. The file should be named authorized_keys, not authorized_key.
The mv command will rename the file and fix the issue. The other options are incorrect because they either do not affect the file name (restorecon or chmod) or do not restart the SSH service (systemct1). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 410.

NEW QUESTION: 257
A Linux administrator cloned an existing Linux server and built a new server from that clone. The administrator encountered the following error after booting the cloned server:

```
Device mismatch detected
```

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

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

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

A. mount disk by device-id

B. fsck -A

C. mount disk by-label

D. mount disk by-blkid

Answer: (SHOW ANSWER)

Explanation

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

NEW QUESTION: 258

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

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


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

A. chgrp -R 755 data/

B. chmod -R 777 data/

C. chattr -R -i data/

D. chown -R data/

Answer: C (LEAVE A REPLY)

Explanation

The command that can be used to resolve the issue of being unable to remove a particular data folder is chattr -R -i data/. This command will use the chattr utility to change file attributes on a Linux file system. The -R option means that chattr will recursively change attributes of directories and their contents. The -i option means that chattr will remove (unset) the immutable attribute from files or directories. When a file or directory has the immutable attribute set, it cannot be modified, deleted, or renamed.

The other options are not correct commands for resolving this issue. The chgrp -R 755 data/ command will change the group ownership of data/ and its contents recursively to 755, which is not a valid group name. The chgrp command is used to change group ownership of files or directories. The chmod -R 777 data/ command will change the file mode bits of data/ and its contents recursively to 777, which means that everyone can read, write, and execute them. However, this will not remove the immutable attribute, which prevents deletion or modification regardless of permissions. The chmod command is used to change file mode bits of files or directories. The chown -R data/ command is incomplete and will produce an error. The chown command is used to change the user and/or group ownership of files or directories, but it requires at least one argument besides the file name. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 7:

Managing Disk Storage; chattr(1) - Linux manual page; chgrp(1) - Linux manual page; chmod(1) - Linux manual page; chown(1) - Linux manual page

NEW QUESTION: 259

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

A. docker run -ti app /bin/sh

B. podman exec -ti app /bin/sh

C. podman run -d app /bin/bash

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

NEW QUESTION: 260
An administrator recently updated the BIND software package and would like to review the default configuration that shipped with this version. Which of the following files should the administrator review?
A. /etc/named.conf.rpmnew
B. /etc/named.conf.rpmsave
C. /etc/named.conf
D. /etc/bind/bind.conf
Answer: A (LEAVE A REPLY)
After installing a new version of a package that includes a configuration file that already exists on the system, such as /etc/httpd/conf/httpd.conf, RPM will create a new file with the .rpmnew extension instead of overwriting the existing file. This allows the administrator to review the default configuration that shipped with this version and compare it with the current configuration before deciding whether to merge or replace the files. The /etc/named.conf.rpmsave file is created by RPM when a package is uninstalled and it contains a configuration file that was modified by the administrator. This allows the administrator to restore the configuration file if needed.
The /etc/named.conf file is the main configuration file for the BIND name server, not the httpd web server. The /etc/bind/bind.conf file does not exist by default in Linux systems. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 19: Managing Packages and Software, page 561.

NEW QUESTION: 261
Which of the following commands is used to configure the default permissions for new files?
A. setenforce
B. sudo
C. umask
D. chmod
Answer: C (LEAVE A REPLY)

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

NEW QUESTION: 262
Based on an organization's new cybersecurity policies, an administrator has been instructed to ensure that, by default, all new users and groups that are created fall within the specified values below.

```
# Min/max values for automatic uid selection in useradd
#
UID_MIN 1000
UID_MAX 60000
# Min/max values for automatic gid selection in groupadd
#
GID_MIN 1000
GID_MAX 60000
```


To which of the following configuration files will the required changes need to be made?
A. /etc/login.defs
B. /etc/security/limits.conf
C. /etc/default/useradd
D. /etc/profile
Answer: A (LEAVE A REPLY)
Explanation
The required changes need to be made to the /etc/login.defs configuration file. The /etc/login.defs file defines the default values for user and group IDs, passwords, shells, and other parameters for user and group creation.
The file contains the directives UID_MIN, UID_MAX, GID_MIN, and GID_MAX, which set the minimum and maximum values for automatic user and group ID selection. The administrator can edit this file and change the values to match the organization's new cybersecurity policies. This is the correct file to modify to accomplish the task. The other options are incorrect because they either do not affect the user and group IDs (/etc/security/limits.conf or /etc/profile) or do not set the default values (/etc/default/useradd). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 15:

Managing Users and Groups, page 463.

NEW QUESTION: 263

Joe, a user, is unable to log in to the Linux system. Given the following output:

```
# grep joe /etc/passwd /etc/shadow
/etc/passwd:joe:x:1001:1001::/home/joe:/bin/nologin
/etc/shadow:joe:$6$3uOw6qWx9876jGhgKJsdfH987634534voj.:18883:0:99999:7:::
```

Which of the following commands would resolve the issue?

A. usermod -s /bin/bash joe

B. pam_tally2 -u joe -r

C. passwd -u joe

D. chage -E 90 joe

Answer: B (LEAVE A REPLY)

Explanation

The command pam_tally2 -u joe -r will resolve the issue of Joe being unable to log in to the Linux system.

The pam_tally2 command is a tool for managing the login counter for the PAM (Pluggable Authentication Modules) system. PAM is a framework for managing authentication and authorization on Linux systems.

PAM allows the administrator to define the rules and policies for accessing various system resources and services, such as login, sudo, ssh, or cron. PAM also supports different types of authentication methods, such as passwords, tokens, biometrics, or smart cards. PAM can be used to implement login restrictions, such as limiting the number of failed login attempts, locking the account after a certain number of failures, or enforcing a minimum or maximum time between login attempts. The pam_tally2 command can display, reset, or unlock the login counter for the users or hosts. The -u joe option specifies the user name that the command should apply to. The -r option resets the login counter for the user. The command pam_tally2 -u joe -r will reset the login counter for Joe, which will unlock his account and allow him to log in to the Linux system. This will resolve the issue of Joe being unable to log in to the Linux system. This is the correct command to use to resolve the issue. The other options are incorrect because they either do not unlock the account (usermod -s /bin/bash joe or passwd -u joe) or do not affect the login counter (chage -E 90 joe). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17: Implementing Basic Security, page 517.

NEW QUESTION: 264

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

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

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

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


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

A. Update the name server in resolv. conf to use an external DNS server.
B. Remove the entry for portal . comptia.org from the local hosts file.
C. Add a network route from the 10.10.10.0/24 to the 192.168.0.0/16.
D. Clear the local DNS cache on the workstation and rerun the host command.
Answer: B (LEAVE A REPLY)
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: 265
As part of the requirements for installing a new application, the swappiness parameter needs to be changed to O. This change needs to persist across re-boots and be applied immediately. A Linux systems administrator is performing this change. Which of the following steps should the administrator complete to accomplish this task?
A. echo "vm. swappiness-()" >> /etc/sysctl . conf && sysctl -p
B. echo "vrn. >> / proc/meminfo && sysctl -a
C. sysctl -v >> / proc/meminfo & & echo "vm. swapiness=0"
D. sysctl -h "vm. swapiness-O" && echo / etc/vmswapiness
Answer: (SHOW ANSWER)
Explanation
To change the swappiness parameter to 0 and make it persistent across reboots and applied immediately, the administrator can perform the following steps:
Append the line vm.swappiness=0 to the file /etc/sysctl.conf using echo "vm.swappiness=0" >> /etc/sysctl.conf (A). This will set the swappiness parameter to 0 for future boots.
Reload the sysctl configuration using sysctl -p (A). This will apply the changes to the current system without rebooting. The other commands will not achieve this task, but either write to a wrong file, use a wrong option, or have a syntax error.
[CompTIA Linux+ Study Guide], Chapter 8: Optimizing Linux Performance, Section: Tuning Kernel Parameters with sysctl
[How to Change Swappiness in Linux]

NEW QUESTION: 266
A systems administrator received a notification that a system is performing slowly. When running the top command, the systems administrator can see the following values:

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

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

A. vmstat

B. strace

C. htop

D. lsof

Answer: (SHOW ANSWER)

Explanation

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

NEW QUESTION: 267

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

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

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

A. restorecon -rv .ssh/authorized_key

B. mv .ssh/authorized_key .ssh/authorized_keys

C. systemct1 restart sshd.service

D. chmod 600 mv .ssh/authorized_key

Answer: (SHOW ANSWER)

Explanation

The command mv .ssh/authorized_key .ssh/authorized_keys will resolve the issue and allow an SSH connection to be established without a password. The issue is caused by the incorrect file name of the authorized key file on the server. The file should be named authorized_keys, not authorized_key.

The mv command will rename the file and fix the issue. The other options are incorrect because they either do not affect the file name (restorecon or chmod) or do not restart the SSH service (systemct1). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 13: Managing Network Services, page 410.

NEW QUESTION: 268

A Linux administrator has been tasked with installing the most recent versions of packages on a RPM-based OS. Which of the following commands will accomplish this task?

A. rpm -a

B. dnf update

C. yum check-update

D. apt-get upgrade

E. yum updateinfo

Answer: B (LEAVE A REPLY)

NEW QUESTION: 269

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 pull https://git.company.com/admin/project.git

C. git checkout https://git.company.com/admin/project.git

D. git branch https://git.company.com/admin/project.git

Answer: A (LEAVE A REPLY)

NEW QUESTION: 270

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

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

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

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

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


Which of the following commands will fix the issue?

A. semanage port -a -t ssh_port_t -p tcp 2222

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

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

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

Answer: A (LEAVE A REPLY)

Explanation

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

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

The other options are incorrect because:

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

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

C: iptables -A INPUT -p tcp --dport 2222 -j ACCEPT

This command will add a rule to the iptables firewall to accept incoming TCP connections on port 2222.

However, this is not enough to fix the issue, as SELinux will still block the SSH server from binding to that port. Moreover, iptables may not be the default firewall service on some Linux distributions, such as Fedora or CentOS, which use firewalld instead.
D: firewall-cmd --zone=public --add-port=2222/tcp
This command will add a rule to the firewalld firewall to allow incoming TCP connections on port 2222 in the public zone. However, this is not enough to fix the issue, as SELinux will still block the SSH server from binding to that port. Moreover, firewalld may not be installed or enabled on some Linux distributions, such as Ubuntu or Debian, which use iptables instead.
References:
How to configure SSH to use a non-standard port with SELinux set to enforcing Change SSH Port on CentOS/RHEL/Fedora With SELinux Enforcing How to change SSH port when SELinux policy is enabled

NEW QUESTION: 271
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 (LEAVE A REPLY)
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: 272

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

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

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

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

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

Answer: D (LEAVE A REPLY)

The list of commands in option D is the correct way to achieve the goal. The commands are as follows:

fallocate -l 1G /ops/app.img creates a 1GB file named app.img under the /ops directory.

mkfs.xfs /ops/app.img formats the file as an XFS filesystem.

mount -o loop /ops/app.img /ops/app mounts the file as a loop device under the /ops/app directory. The other options are incorrect because they either use the wrong commands (dd or truncate instead of fallocate), the wrong options (-t or -f instead of -o), or the wrong order of arguments (/ops/app.img /ops/app instead of /ops/app /ops/app.img). Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 10: Managing Storage, pages 323-324.