Updated Jun-2024 100% Cover Real XK0-005 Exam Questions Make Sure You 100% Pass [Q77-Q99

Updated Jun-2024 100% Cover Real XK0-005 Exam Questions Make Sure You 100% Pass XK0-005 dumps Accurate Questions and Answers with Free and Fast Updates QUESTION 77

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:

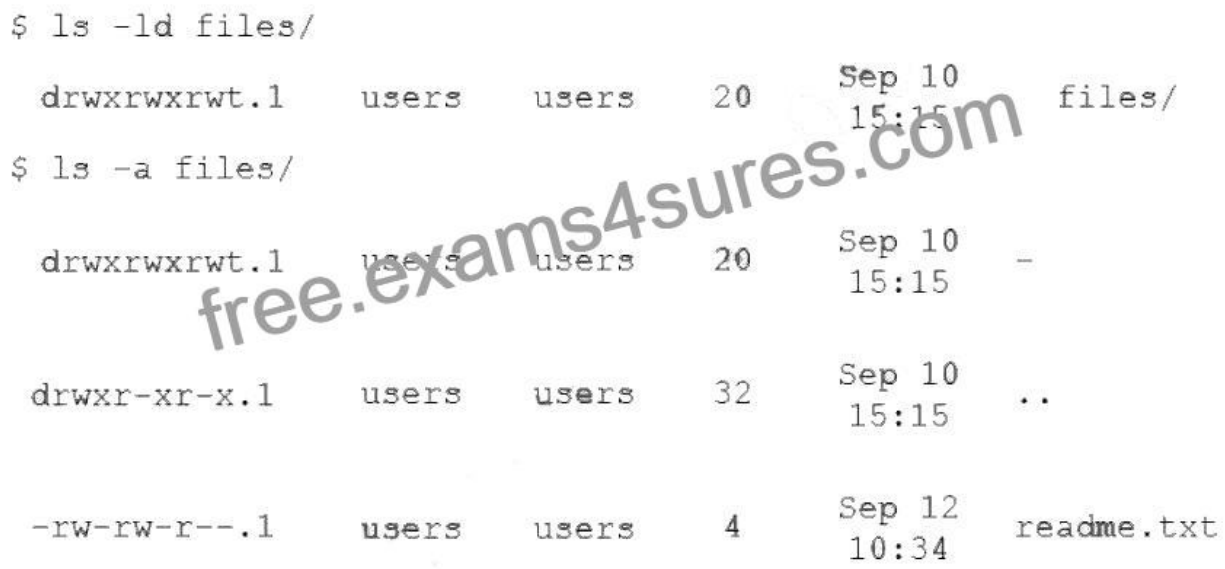

Which of the following commands should the administrator run NEXT to allow the file to be renamed by any user?

* chgrp reet files
* chacl -R 644 files
* chown users files
* chmod -t files

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. Reference: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 8: Managing Users and Groups; chmod(1) – Linux manual page

QUESTION 78

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.

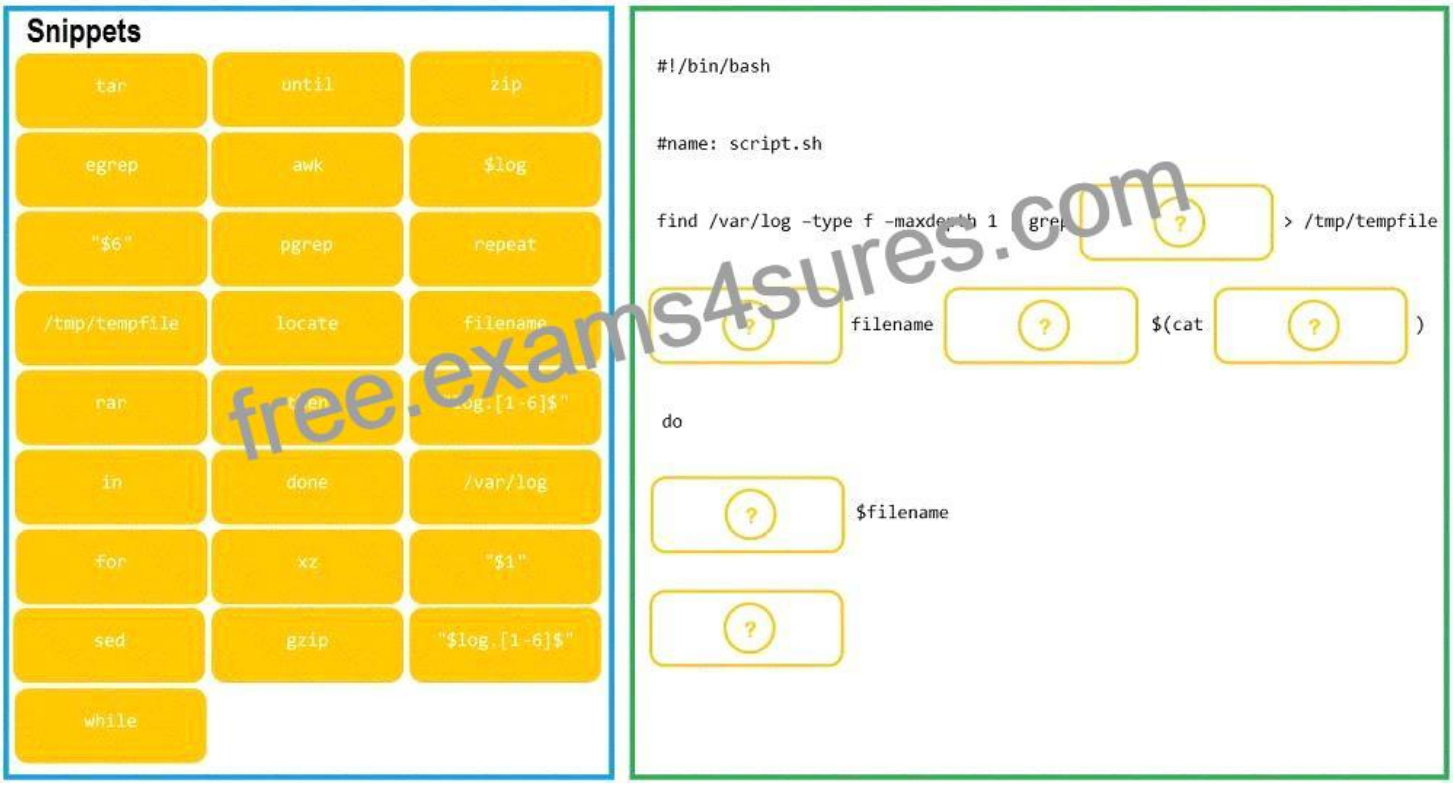

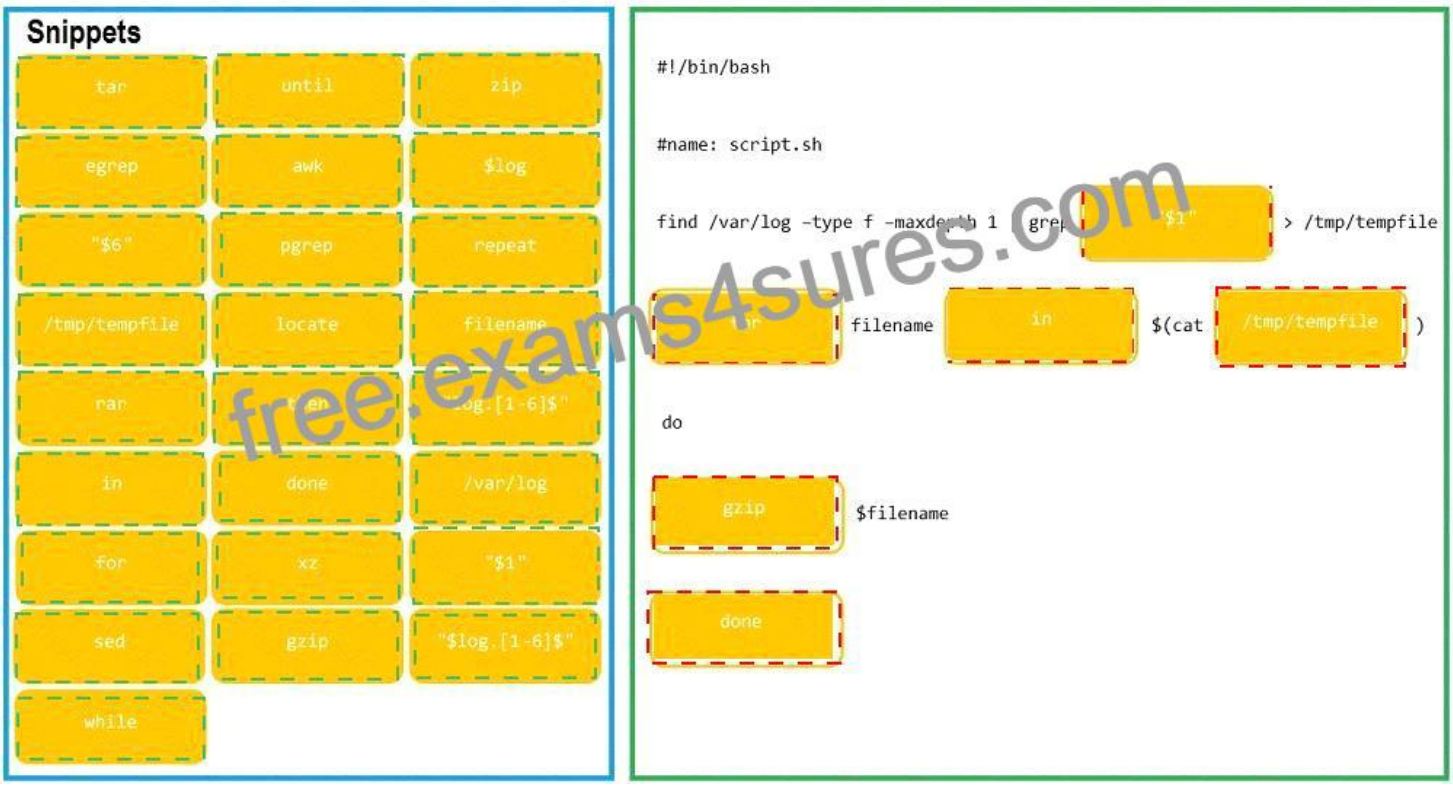

QUESTION 79

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.

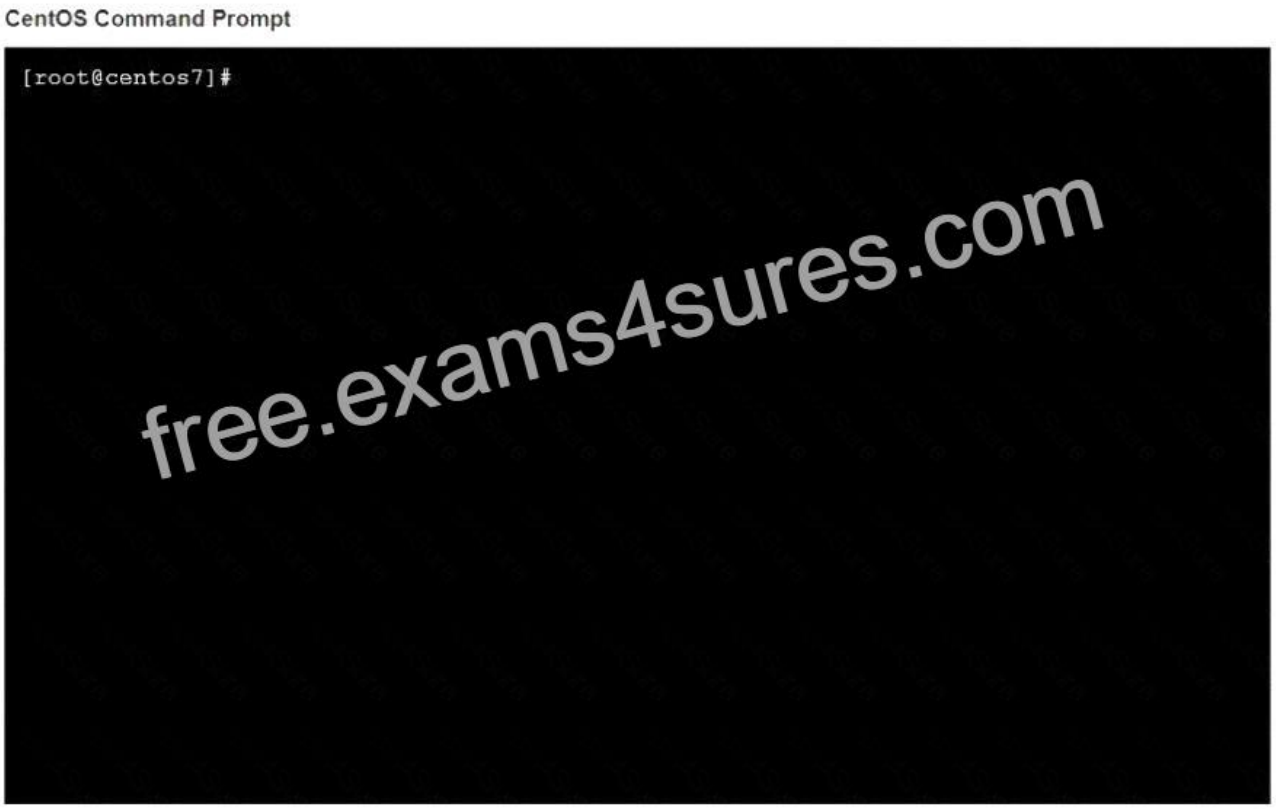

See the explanation below.

Explanation

yum install httpd

systemctl –now enable httpd

systemctl status httpd

netstat -tunlp | grep 80

pkill <processname>

systemctl restart httpd

systemctl status httpd

QUESTION 80

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?

* lsblk
* fdisk
* df -h

* du -ah

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. Reference: [CompTIA Linux+ (XK0-005) Certification Study Guide], Chapter 14: Managing Disk Storage, page 454.

QUESTION 81

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

* systemctl mask httpd
* systemctl disable httpd
* systemctl stop httpd
* systemctl reload httpd

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

C) systemctl stop httpd will only stop the service if it is currently running, but it will not prevent it from being started again.

D) systemctl reload httpd will only reload the configuration files of the service, but it will not stop or disable it.

QUESTION 82

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

* kill -1
* kill -9
* kill -15
* kill -HUP
* kill -TERM

QUESTION 83

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?
* The IP address 0.0.0.0 is not valid.
* The application is listening on the loopback interface.
* The application is listening on port 1234.
* The application is not running.

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.

QUESTION 84

A User on a Linux workstation needs to remotely start an application on a Linux server and then forward the graphical display of that application back to the Linux workstation. Which of the following would enable the user to perform this action?
* ssh -X user@server application
* ssh -y user@server application
* ssh user@server application
* ssh -D user@server application

The ssh -X option enables X11 forwarding, which allows the user to run graphical applications on the remote server and display them on the local workstation. The user needs to specify the username, the server address, and the application name after the ssh -X command. The remote server also needs to have X11Forwarding enabled and xauth installed for this to work.

Reference:

The web search result 8 explains how to run a GUI application through SSH by configuring both the SSH client and server.

The web search result 6 provides a detailed answer on how to forward X over SSH to run graphics applications remotely, with examples and troubleshooting tips.

The CompTIA Linux+ Certification Exam Objectives mention that the candidate should be able to “use SSH for remote access and management” as part of the System Operation and Maintenance domain1.

QUESTION 85

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.)

* Add #!/bin/bash to 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.
* Add #!//bin/bash to the top of the script.
* Restart the computer to enable the new service.
* Create a unit file for the new service in /etc/init.d with the name helpme.service in the location.
* Shut down the computer to enable the new service.

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.

QUESTION 86

A Linux administrator booted up the server and was presented with a non-GUI terminal. The administrator ran the command systemct1 isolate graphical.target and rebooted the system by running systemct1 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?

* The administrator did not reboot the server properly.
* The administrator did not set the default target to basic.target.
* The administrator did not set the default target to graphical.target.
* The administrator did not shut down the server properly.

Explanation

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 systemct1 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 systemct1 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 systemct1 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. References: systemct1(1) – Linux manual page; How to Change Runlevels (targets) in SystemD

QUESTION 87

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?

* mount disk by device-id
* fsck -A
* mount disk by-label
* mount disk by-blkid

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.

QUESTION 88

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?

* dig @example.com 10.10.10.20 a
* dig @10.10.10.20 example.com mx
* dig @example.com 10.10.10.20 ptr
* dig @10.10.10.20 example.com ns

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.

QUESTION 89

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?

* The DNS address has been commented out in the configuration file.
* The search entry in the /etc/resolv.conf file is incorrect.
* Wired connection 1 is offline.
* No default route is defined.

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

QUESTION 90

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?
* The IP address 0.0.0.0 is not valid.
* The application is listening on the loopback interface.
* The application is listening on port 1234.
* The application is not running.

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.

QUESTION 91

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

QUESTION 92

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?
* chmod 775
* umask. 002
* chactr -Rv
* chown -cf

QUESTION 93

An administrator needs to increase the system priority of a process with PID 2274. Which of the following commands should the administrator use to accomplish this task?
* renice -n -15 2274
* nice -15 2274
* echo “-15” > /proc/PID/2274/priority
* ps -ef I grep 2274

Explanation

The renice command is used to change the priority of a running process by specifying its PID and the new nice value. The -n flag indicates the amount of change in the nice value, which can be positive or negative. A lower nice value means a higher priority, so -15 will increase the priority of the process with PID 2274. The administrator needs to have root privileges to do this.

References:

The renice command is listed as one of the commands to manipulate process priority in the web search result 1.

The renice command is also explained with examples in the web search result 2.

The CompTIA Linux+ Certification Exam Objectives mention that the candidate should be able to

“manage process execution priorities” as part of the System Operation and Maintenance domain1.

QUESTION 94

During a security scan, the password of an SSH key file appeared to be too weak and was cracked. Which of the following commands would allow a user to choose a stronger password and set it on the existing SSH key file?
* passwd
* ssh
* ssh-keygen
* pwgen

Explanation

The command that would allow a user to choose a stronger password and set it on the existing SSH key file is ssh-keygen -p -f <keyfile>. This command uses the ssh-keygen tool, which is used to generate, manage, and convert authentication keys for SSH. The -p option stands for passphrase, and it allows the user to change or remove the passphrase of an existing private key file. The -f option specifies the filename of the key file. The command will prompt the user for the old passphrase, and then for the new passphrase twice.

The other options are not correct commands for changing the password of an SSH key file. The passwd command is used to change the password of a user account on a Linux system, not an SSH key file. The ssh command is used to log in to a remote system using SSH, not to change the password of an SSH key file. The pwgen command is used to generate random passwords, not to change the password of an SSH key file.

References: ssh-keygen(1) – Linux manual page; How To: Change Passphrase for SSH Private Key – Unix Tutorial

QUESTION 95

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?

* dig @example.com 10.10.10.20 a
* dig @10.10.10.20 example.com mx
* dig @example.com 10.10.10.20 ptr
* dig @10.10.10.20 example.com ns

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.

QUESTION 96

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

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


The administrator notices the file . DS STORE should not be included and deletes it from the online repository. Which of the following should the administrator run from the root of the local repository before the next commit to ensure the file is not uploaded again in future commits?
* rm -f .DS STORE && git push
* git fetch && git checkout .DS STORE
* rm -f .DS STORE && git rebase origin main
* echo .DS STORE >> .gitignore

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.

QUESTION 97

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?
* git fetch
* git checkout
* git clone
* git branch

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

QUESTION 98

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

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

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

Given this scenario, which of the following should the administrator do to recover this volume?
* Reboot the server. The volume will automatically go back to linear mode.
* Replace the failed drive and reconfigure the mirror.
* Reboot the server. The volume will revert to stripe mode.
* Recreate the logical volume.

QUESTION 99

A systems administrator is tasked with setting up key-based SSH authentication. In which of the following locations should the administrator place the public keys for the server?
* ~/.sshd/authkeys
* ~/.ssh/keys
* ~/.ssh/authorized_keys

* ~/.ssh/keyauth

Real XK0-005 Quesions Pass Certification Exams Easily:
https://www.exams4sures.com/CompTIA/XK0-005-practice-exam-dumps.html]