XK0-005 Q&As

CompTIA Linux+

Pass CompTIA XK0-005 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4itsure.com/xk0-005.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by CompTIA Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

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

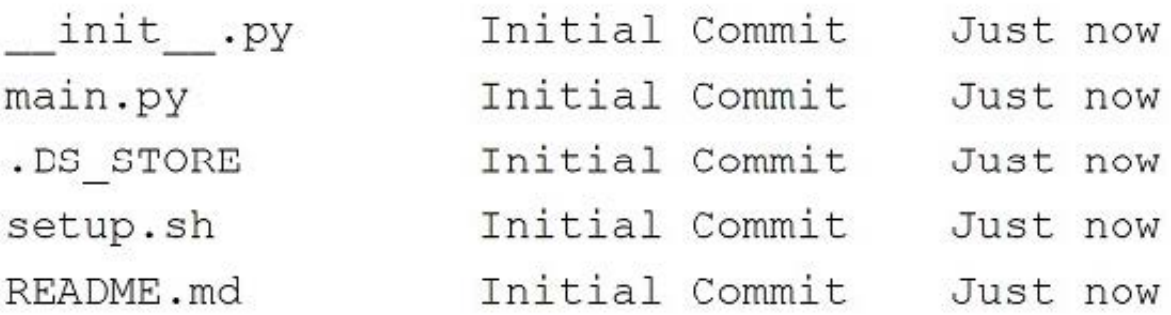

The administrator notices the file . DS STORE should not be included and deletes it from the online repository. Which of the following should the administrator run from the root of the local repository before the next commit to ensure the file is not uploaded again in future commits?

A. rm -f .DS STORE andand git push

B. git fetch andand git checkout .DS STORE

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

D. echo .DS STORE >> .gitignore

Correct Answer: D

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 andand 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 andand 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 andand 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 2

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

A. Docker

B. On-premises systems

C. Cloud-based systems

D. Kubernetes

Correct Answer: D

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

QUESTION 3

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.

Correct Answer: D

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

QUESTION 4

A systems administrator made some changes in the ~/.bashrc file and added an alias command. When the administrator tried to use the alias command, it did not work. Which of the following should be executed FIRST?

A. source ~/.bashrc

B. read ~/.bashrc

C. touch ~/.bashrc

D. echo ~/.bashrc

Correct Answer: A

Explanation: The command source ~/.bashrc should be executed first to use the alias command. The source command reads and executes commands from a file in the current shell environment. The ~/.bashrc file is a configuration file that contains commands and aliases that are executed when a new bash shell is started. The administrator made some changes in the ~/.bashrc file and added an alias command, but the changes are not effective until the file is sourced or a new shell is started. The command source ~/.bashrc will reload the file and make the alias command available. The

other options are incorrect because they either do not execute the commands in the file (read, touch, or echo) or do not affect the current shell environment (read or echo). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter

9: Working with the Linux Shell, page 295.

QUESTION 5

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

A. route -e get to 192.168.1.40 from 10.0.2.15

B. ip route get 192.163.1.40 from 10.0.2.15

C. ip route 192.169.1.40 to 10.0.2.15

D. route -n 192.168.1.40 from 10.0.2.15

Correct Answer: B

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

QUESTION 6

A systems administrator needs to reconfigure a Linux server to allow persistent IPv4 packet forwarding. Which of the following commands is the correct way to accomplish this task?

A. echo 1 > /proc/sys/net/ipv4/ipv_forward

B. sysctl -w net.ipv4.ip_forward=1

C. firewall-cmd --enable ipv4_forwarding

D. systemct1 start ipv4_forwarding

Correct Answer: B

Explanation: The command sysctl -w net.ipv4.ip_forward=1 enables IPv4 packet forwarding temporarily by setting the kernel parameter net.ipv4.ip_forward to 1. To make this change persistent, the administrator needs to edit the file /etc/sysctl.conf and add the line net.ipv4.ip_forward = 1. The other options are incorrect because they either use the wrong file (/proc/sys/net/ipv4/ipv_forward), the wrong command (firewall- cmd or systemct1), or the wrong option (--enable or start). References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 12: Managing Network Connections, page 378.

QUESTION 7

A cloud engineer wants to delete all unused networks that are not referenced by any container. Which of the following commands will achieve this goal?

A. docker network erase

B. docker network clear

C. docker network prune

D. docker network rm

Correct Answer: C

The docker command is used to manage Docker containers, images, networks, volumes, and other resources on a Linux system. Docker is a platform that allows users to run applications in isolated environments called containers. Docker also provides networking features that allow users to create and manage networks for containers. To delete all unused networks that are not referenced by any container, the cloud engineer can use the docker network prune command. This command will remove all networks that have no containers connected to them. The statement C is correct. The statements A, B, and D are incorrect because they do not delete all unused networks. The docker network erase and docker network clear commands do not exist. The docker network rm command deletes a specific network by name or ID, but not all unused networks. References: [How to Manage Docker Networks]

QUESTION 8

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. systemct1 status

B. systemct1 stop

C. systemct1 reinstall

D. systemct1 daemon-reload

Correct Answer: D

Explanation: After installing a new version of a package that includes a new version of the corresponding service file, the systemct1 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 systemct1 status command will

display information about a service unit, but it will not reload the configuration. The systemct1 stop command will stop a service unit, but it will not reload the configuration. The systemct1 reinstall command does not exist. References:

CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 17:

System Maintenance and Operation, page 518.

QUESTION 9

An administrator transferred a key for SSH authentication to a home directory on a remote server. The key file was moved to .ssh/authorized_keys location in order to establish SSH connection without a password. However, the SSH

command still asked for the password.

Given the following output:

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

Which of the following commands would resolve the issue?

A. restorecon .ssh/authorized_keys

B. ssh_keygen -t rsa -o .ssh/authorized_keys

C. chown root:root .ssh/authorized_keys

D. chmod 600 .ssh/authorized_keys

Correct Answer: D

Explanation: The command that would resolve the issue is chmod 600 .ssh/authorized_keys. This command will change the permissions of the .ssh/authorized_keys file to 600, which means that only the owner of the file can read and write it. This is necessary for SSH key authentication to work properly, as SSH will refuse to use a key file that is accessible by other users or groups for security reasons. The output of ls -l shows that currently the .ssh/authorized_keys file has permissions of 664, which means that both the owner and group can read and write it, and others can read it. The other options are not correct commands for resolving the issue. The restorecon .ssh/authorized_keys command will restore the default SELinux security context for the .ssh/authorized_keys file, but this will not change its permissions or ownership. The ssh_keygen -t rsa -o .ssh/authorized_keys command is invalid because ssh_keygen is not a valid command (the correct command is ssh-keygen), and the -o option is used to specify a new output format for the key file, not the output file name. The chown root:root .ssh/authorized_keys command will change the owner and group of the .ssh/ authorized_keys file to root, but this will not change its permissions or make it accessible by the user who wants to log in with SSH key authentication. References: How to Use Public Key Authentication with SSH; chmod(1) - Linux manual page

QUESTION 10

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.

Correct Answer: D

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 11

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

Correct Answer: D

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.

QUESTION 12

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

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

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

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

A. Reboot the server. The volume will automatically go back to linear mode.

B. Replace the failed drive and reconfigure the mirror.

C. Reboot the server. The volume will revert to stripe mode.

D. Recreate the logical volume.

Correct Answer: B

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

QUESTION 13

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

Correct Answer: A

Explanation: The command that will accomplish the task of running a container with the httpd server inside and exposing it at port 443 of the Linux host machine while it internally listens on port 8443 is podman run -d -p 443:8443

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

QUESTION 14

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

Correct Answer: C

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:

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.

QUESTION 15

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

```
# 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
```

A. chgrp -R 755 data/

B. chmod -R 777 data/

C. chattr -R -i data/

D. chown -R data/

Correct Answer: C

Explanation: The command that can be used to resolve the issue of being unable to remove a particular data folder is chattr -R -i data/. This command will use the chattr utility to change file attributes on a Linux file system. The -R option means that chattr will recursively change attributes of directories and their contents. The -i option means that chattr will remove (unset) the immutable attribute from files or directories. When a file or directory has the immutable attribute set, it cannot be modified, deleted, or renamed. The other options are not correct commands for resolving this issue. The chgrp -R 755 data/ command will change the group ownership of data/ and its contents recursively to 755, which is not a valid group name. The chgrp command is used to change group ownership of files or directories. The chmod -R 777 data/ command will change the file mode bits of data/ and its contents recursively to 777, which means that everyone can read, write, and execute them. However, this will not remove the immutable attribute, which prevents deletion or modification regardless of permissions. The chmod command is used to change file mode bits of files or directories. The chown -R data/ command is incomplete and will produce an error. The chown command is used to change the user and/or group ownership of files or directories, but it requires at least one argument besides the file name. References: CompTIA Linux+ (XK0-005) Certification Study Guide, Chapter 7: Managing Disk Storage; chattr(1) - Linux manual page; chgrp(1) - Linux manual page; chmod(1) - Linux manual page; chown(1) - Linux manual page