Comptia

(XK0-005)

CompTIA Linux+

Total: **275 Questions**

Link:

Question: 1

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

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

Answer: A

Explanation:

The correct approach to identifying the correct kernel version after accidentally deleting /boot/vmlinuz before a reboot is to use commands that reveal installed kernel packages and the currently running kernel.

Option A, rpm -qa | grep kernel; uname -a, achieves this. rpm -qa | grep kernel lists all installed RPM packages and filters the output to show only those containing "kernel" in their name. This helps identify the installed kernel versions. uname -a displays comprehensive information about the current kernel, including its version and build details. By comparing the output of both commands, the administrator can pinpoint the correct vmlinuz file to restore or reinstall.

Option B, yum -y update; shutdown -r now, attempts to update all packages and then reboot the server. This is risky because the server might fail to boot after the update if the kernel isn't properly replaced first. Furthermore, it does not help in identifying the correct kernel to use immediately. Updating without first knowing the current kernel version could lead to incompatibility issues if the wrong kernel image gets installed.

Option C, cat /etc/centos-release; rpm -Uvh --nodeps, displays the CentOS version and attempts to upgrade an RPM package without dependency checking. Knowing the CentOS version might be helpful contextually, but doesn't directly provide the installed kernel versions. Ignoring dependencies using --nodeps can cause system instability and is generally discouraged.

Option D, telinit 1; restorecon -Rv /boot, switches the system to single-user mode and restores SELinux contexts to the /boot directory. While fixing SELinux contexts is crucial for security, switching to single user mode would likely require a reboot, and is not helpful in this context since you wish to avoid rebooting the server if possible. restorecon only addresses SELinux attributes, not the missing vmlinuz file. It is also a dangerous command since switching to single user mode would likely result in a reboot given the missing vmlinuz file.

Therefore, Option A offers a way to identify the correct kernel version without necessitating a risky update or a potentially failed boot sequence, making it the most suitable option. It aligns with best practices for system administration by providing information before taking potentially disruptive actions.

Further reading:

rpm Command:https://www.redhat.com/sysadmin/linux-package-management
uname Command:https://man7.org/linux/man-pages/man1/uname.1.html

Question: 2

A cloud engineer needs to change the secure remote login port from 22 to 49000. Which of the following files should the engineer modify to change the port number to the desired value?

A./etc/host.conf
B./etc/hostname
C./etc/services
D./etc/ssh/sshd_config

Answer: D

Explanation:

The correct answer is **D. /etc/ssh/sshd_config**. This file is the primary configuration file for the SSH daemon (sshd) on Linux systems, which handles secure remote logins using the SSH protocol. Modifying this file allows administrators to customize various aspects of the SSH server's behavior, including the port on which it listens for incoming connections.

To change the SSH port, the engineer needs to edit the Port directive within the /etc/ssh/sshd_config file. By default, this directive is usually set to 22. Changing it to 49000 instructs the SSH daemon to listen for connections on that port instead. After modifying the file, the SSH daemon must be restarted or reloaded for the changes to take effect. This ensures that the server starts using the newly configured port. It's also important to update firewall rules to allow traffic on the new port, preventing connection issues.

The other options are incorrect:

/etc/host.conf is used for hostname resolution and is not related to SSH configuration.
/etc/hostname specifies the system's hostname, not SSH settings.
/etc/services maps service names to port numbers and protocols but doesn't configure individual services like SSH. While modifying it might seem relevant, the SSH daemon specifically reads its configuration from /etc/ssh/sshd_config. Modifying /etc/services will not change the port SSH listens on.

Therefore, /etc/ssh/sshd_config is the only file that directly controls the SSH daemon's port configuration, making it the correct answer.

Authoritative Links for further research:

SSH Configuration Options:https://man.openbsd.org/sshd_config **SSH - ArchWiki:**https://wiki.archlinux.org/title/SSH

Question: 3

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 reflog
B.git pull
C.git status
D.git push

Answer: B

Explanation:

The correct command is git pull. This command is designed specifically to synchronize a local repository with a remote repository. It effectively combines two operations: git fetch and git merge. First, git fetch retrieves the latest changes (commits, branches, tags) from the remote repository but without automatically integrating them into your local working directory. Then, git merge integrates the fetched changes into the current branch

of your local repository. This ensures your local copy reflects the most up-to-date state of the main repository.

Option A, git reflog, is used to view a log of changes made to the local repository's reference pointers (like HEAD, branch tips), allowing you to recover from accidental branch resets or deletions, but doesn't synchronize with a remote. Option C, git status, displays the state of the working directory and staging area, showing which files have been modified, staged, or are untracked, but it doesn't retrieve updates from the remote. Option D, git push, is used to upload local repository content to a remote repository, the opposite of what is required here. You use git push to send your local changes to the main repository. Using git pull is fundamental to maintaining synchronization in a collaborative Git workflow, a cornerstone of modern software development and cloud-based version control systems.

For further information:

Git pull documentation: https://git-scm.com/docs/git-pull
Git basics: https://git-scm.com/book/en/v2/Git-Basics-Getting-a-Git-Repository

Question: 4

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

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

Answer: D

Explanation:

The correct command to redirect HTTP traffic to a new proxy server using iptables involves modifying the NAT (Network Address Translation) table. Specifically, we need to alter the destination of incoming HTTP requests.

Option D, iptables -t nat -A PREROUTING -p tcp --dport 80 -j DNAT --to-destination 192.0.2.25:3128, is the correct choice because it accurately reflects the steps required for this redirection. Let's break down why:

-t nat: Specifies that we are working with the NAT table, which is responsible for network address translation. This is essential because we are changing the destination IP and port.
-A PREROUTING: Appends the rule to the PREROUTING chain. The PREROUTING chain is the first chain that incoming packets traverse before routing decisions are made. This is where we want to intercept HTTP traffic and modify its destination before it reaches its original intended destination.
-p tcp: Specifies that the rule applies only to TCP traffic, which is the protocol used by HTTP.
--dport 80: Specifies that the rule applies only to traffic destined for port 80, which is the standard port for HTTP traffic.
-j DNAT: Specifies the target of the rule as DNAT (Destination NAT). DNAT is used to change the destination IP address and/or port of a packet.
--to-destination 192.0.2.25:3128: Specifies the new destination IP address (192.0.2.25) and port (3128) to which the HTTP traffic should be redirected. This is the IP and port of the proxy server.

Option A is incorrect because it uses -D, which deletes a rule instead of adding one. Option B is incorrect because it redirects traffic destined for port 81 instead of the standard HTTP port 80 and sets incorrect new port 3129. Option C is incorrect as it filters based on source port instead of destination port, which is

inappropriate for the task, and also redirects to incorrect new port 3129. -I also inserts a rule at the beginning, which might cause conflicts if other rules exist.

In summary, option D correctly appends a DNAT rule to the PREROUTING chain within the NAT table to redirect incoming TCP traffic on port 80 to the specified proxy server on port 3128, effectively redirecting all HTTP traffic.

Authoritative Links:

iptables tutorial:https://www.netfilter.org/documentation/
Understanding iptables chains:https://www.digitalocean.com/community/tutorials/understanding-iptables-chains-and-how-they-filter-traffic

Question: 5

Developers have requested implementation of a persistent, static route on the application server. Packets sent over the interface eth0 to 10.0.213.5/32 should be routed via 10.0.5.1. Which of the following commands should the administrator run to achieve this goal?

A.route -i etho -p add 10.0.213.5 10.0.5.1
B.route modify eth0 +ipv4.routes "10.0.213.5/32 10.0.5.1"
C.echo "10.0.213.5 10.0.5.1 eth0" > /proc/net/route
D.ip route add 10.0.213.5/32 via 10.0.5.1 dev eth0

Answer: D

Explanation:

The correct command to add a persistent static route in Linux is ip route add. This command provides the necessary syntax to specify the destination network, gateway, and interface for routing packets. In this scenario, packets destined for 10.0.213.5/32 should be routed through the gateway 10.0.5.1 using the eth0 interface. The ip route add 10.0.213.5/32 via 10.0.5.1 dev eth0 command precisely fulfills these requirements.

Option A is incorrect because the route command does not have a -i option to specify the interface, and -p (persistent) may not work as expected on all Linux distributions without additional configuration. Option B uses an incorrect syntax for modifying routes and lacks a standard Linux command. Option C attempts to write directly to /proc/net/route, which is generally not recommended for persistent configuration, as its contents are dynamically generated by the kernel. The changes made through this method will be lost after reboot. Furthermore, this option does not guarantee persistence across system restarts. To make the route permanent, you typically need to add it to a configuration file specific to your Linux distribution's networking system (e.g., /etc/network/interfaces on Debian-based systems, or network scripts in /etc/sysconfig/network-scripts on Red Hat-based systems). The ip route command coupled with persistence configuration is a more robust approach.

For further research, refer to these resources:

ip command documentation:man ip in your Linux terminal
Red Hat Enterprise Linux documentation on configuring static routes:
https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html/configuring_and_managing_networking/configuring-static-routes_configuring-and-managing-networking
Debian Wiki on NetworkConfiguration:https://wiki.debian.org/NetworkConfiguration

Question: 6

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

Explanation:

"if [-f "$filename"]; then" is a shell script that tests for the existence of a file named "$filename". The syntax of this command is as follows:"if" is a shell construct that allows you to execute a command or a series of commands only if a certain condition is met."[-f "$filename"]" is a test command that returns true if "$filename" is a regular file (i.e., not a directory or a symbolic link). The "-f" option is used to test for the existence of a regular file. The "$filename" variable is enclosed in double quotes to allow for the possibility of spaces in the file name."then" is a keyword that specifies the start of the commands to be executed if the test is true.

Question: 7

DRAG DROP -
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 related files in /var/log without recursion.
INSTRUCTIONS:
Drag and drop snippets to 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.

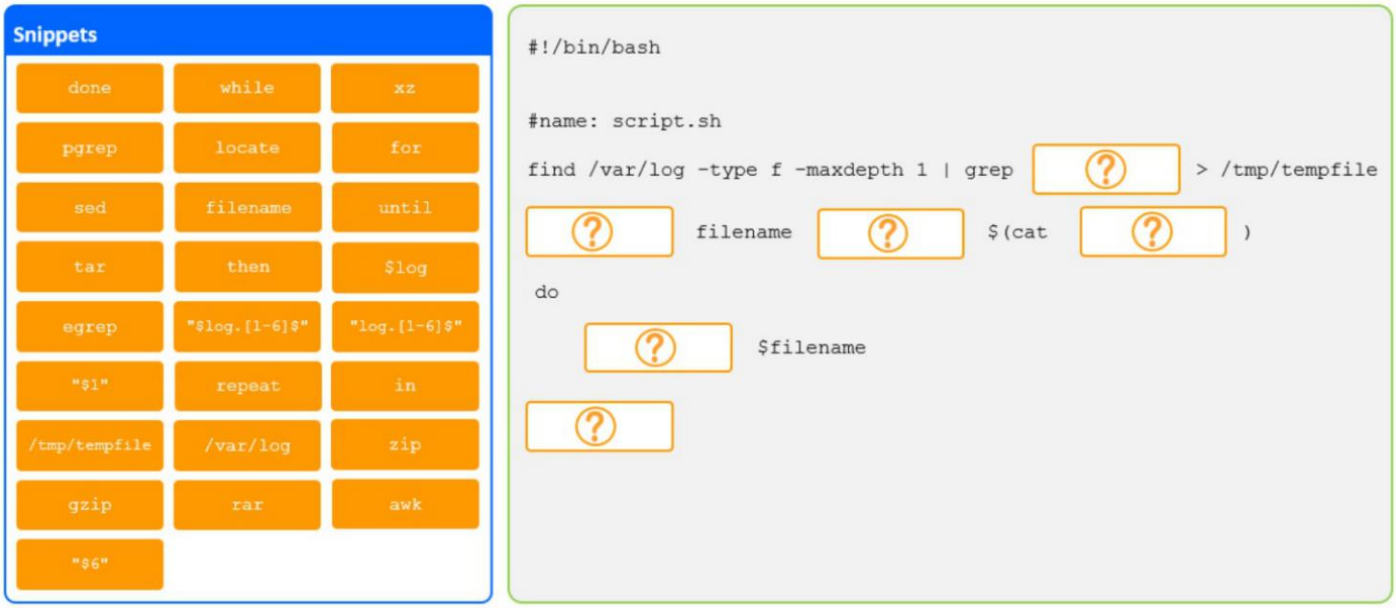

Answer:

Explanation:

#!/bin/bash

#name: script.sh

find /var/log -type f -maxdepth 1 | grep "log.[1-6]$" > /tmp/tempfile for

filename in $(cat /tmp/tempfile); do

gzip $filename

done

Question: 8

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

Explanation:

In this example, Terraform is used to deploy an Amazon Web Services (AWS) EC2 instance. The first line of the code defines the AWS provider, and the region attribute is set to "us-west-2".The second section of the code defines the aws_instance resource. The ami attribute specifies the Amazon Machine Image (AMI) ID to use, and the instance_type attribute specifies the type of instance to deploy. The tags attribute is used to add a tag to the instance, with the key "Name" and the value "example-instance".This code is a simple example of how Terraform can be used to deploy a cloud-based server. In a real-world scenario, the code would likely be more complex and include additional resources and configuration options, such as security groups, subnets, and more.

Question: 9

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

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

Answer: A

Explanation:

LDAP (Lightweight Directory Access Protocol) is a widely used, open-standard protocol for accessing and maintaining distributed directory information services. Its hierarchical, tree-like structure makes it ideal for storing information about network users, groups, devices, and other objects. In a Linux environment, LDAP can serve as a central repository for user and group accounts, streamlining user management across multiple systems. Instead of managing user accounts individually on each Linux server, administrators can manage them centrally through the LDAP server.

When a user attempts to log in to a Linux system configured to use LDAP, the system queries the LDAP server to authenticate the user and retrieve user information like username, password, group memberships, and home directory. This eliminates the need to duplicate user accounts across multiple systems, improving security and simplifying administration. Using LDAP also enables centralized password management, where password policies and changes are applied consistently across all connected systems. This helps enforce strong password policies and reduces the risk of unauthorized access.

The alternative options are less suited for serving as a central user and group repository. MFA (Multi-Factor Authentication) is an authentication method, not a directory service. SSO (Single Sign-On) relies on a central authentication service but doesn't necessarily store user data in a directory format like LDAP does; instead, it focuses on allowing users to authenticate once and gain access to multiple applications. PAM (Pluggable Authentication Modules) is a flexible framework for authentication on a single Linux system but doesn't inherently provide centralized storage across multiple systems; while PAM can integrate with LDAP, PAM itself is not the repository.

Therefore, LDAP is the correct answer because it provides a standardized and efficient way to centralize user and group information in a Linux environment, making user management and security more manageable.

Further Research:

LDAP Overview:https://ldap.com/
OpenLDAP:https://www.openldap.org/
Linux LDAP Configuration:https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html/configuring_authentication_with_ldap/index

Question: 10

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

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

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

Answer: D

Explanation:

D. The network interface cable is not connected to a switch.The output shows that the state of the network interface eth0 is "DOWN". This means that the interface is not currently transmitting or receiving data. The "NO-CARRIER" status also suggests that the interface is not connected to a network.The most probable cause of the issue is that the network interface cable is not connected to a switch, or that the switch the cable is connected to is not functioning correctly. When a network interface is down, it usually means that there is an issue with the physical or logical connection to the network.

Question: 11

A Linux administrator was asked to run a container with the httpd server inside. This container should be exposed at port 443 of a Linux host machine while it internally listens on port 8443. Which of the following commands will accomplish this task?

A.podman run -d -p 443:8443 httpd
B.podman run -d -p 8443:443 httpd
C.podman run –d -e 443:8443 httpd
D.podman exec -p 8443:443 httpd

Answer: A

Explanation:

The correct answer is A: `podman run -d -p 443:8443 httpd`. This command effectively maps the host's port 443 to the container's port 8443, allowing external access on the host's port 443 to be routed to the httpd server running inside the container on port 8443.

Here's a breakdown:

`podman run`: This initiates the creation and running of a new container.
`-d`: This flag runs the container in detached mode, meaning it runs in the background.
`-p 443:8443`: This is the crucial port mapping option. The syntax `host_port:container_port` specifies the mapping. In this case, it maps the host's port 443 (the first number) to the container's port 8443 (the second number). Incoming traffic on the host's port 443 will be forwarded to port 8443 within the container. This satisfies the requirement that the container is exposed on port 443 of the host.
`httpd`: This specifies the image to be used for the container, in this case, an httpd (Apache web server) image.

Option B is incorrect because it reverses the port mapping; it would map the host's port 8443 to the container's port 443, the opposite of what's required. Option C uses `-e`, which is for environment variables, not port mapping. Option D uses `podman exec`, which is for executing a command inside an already running container, not for creating a container and mapping ports.

Port mapping is a core concept in containerization. It allows you to expose services running within a container on specific ports on the host machine. Without port mapping, the services inside the container would only be accessible from within the container itself or other containers on the same network, defeating the purpose of exposing the service to the outside world.

For more information on Podman and port mapping, you can refer to the official Podman documentation:

Podman Documentation:https://docs.podman.io/en/latest/
Red Hat's explanation of Podman's usage:https://www.redhat.com/sysadmin/podman-basics

Question: 12

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

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

Answer: B

Explanation:

The correct answer is **B. podman exec -ti app /bin/sh**. Here's a detailed justification:

The scenario requires accessing a running container to analyze application logs. This means we need to execute a command inside the container's environment. The exec command in both Docker and Podman is designed for this purpose. It allows running a new process (in this case, a shell like /bin/sh or /bin/bash) within an existing container.

Option A, docker run -ti app /bin/sh, is incorrect because docker run creates a new container, potentially based on the 'app' image, but doesn't enter an existing running one. It would be used to start a new container instance, not to interact with one that's already running.

Option C, podman run -d app /bin/bash, is also incorrect for the same reason as option A. Additionally, the -d flag runs the container in detached mode (in the background), which is counterproductive since the goal is to interact with the container.

Option D, docker exec -d app /bin/bash, is incorrect because while docker exec is the correct command structure for Docker, the -d flag again runs the process in detached mode. We specifically need interactive access to the container.

The -ti flags in option B (and potentially used with docker exec if option B were changed to Docker) are crucial. -t allocates a pseudo-TTY (terminal), and -i keeps STDIN open, allowing interactive input from the administrator. This is necessary to type commands within the container's shell after it's opened. podman exec -ti app /bin/sh will execute /bin/sh (a shell) inside the container named 'app', giving the administrator a shell prompt from which to view logs and perform other analysis.

Therefore, podman exec -ti app /bin/sh is the only command that enables interactive access to a running container's shell, allowing the administrator to analyze logs stored within the container. While docker exec -ti app /bin/sh would also work if the question explicitly stated the application was containerized using Docker, option B is more general since Podman is designed as a Docker-compatible container engine.

Relevant Resources:

Podman exec command:https://docs.podman.io/en/latest/markdown/podman-exec.1.html
Docker exec command:https://docs.docker.com/engine/reference/commandline/exec/

Question: 13

A systems administrator needs to clone the partition /dev/sdc1 to /dev/sdd1. Which of the following commands will

accomplish this task?

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

Answer: C

Explanation:

The correct command to clone a partition from /dev/sdc1 to /dev/sdd1 is dd if=/dev/sdc1 of=/dev/sdd1. The dd command is a powerful utility in Linux used for copying and converting data at a low level. dd reads data from an input source (if) and writes it to an output destination (of), block by block. In this context, it directly copies the raw data from the source partition /dev/sdc1 to the destination partition /dev/sdd1, effectively cloning it.

Option A, tar -cvzf /dev/sdd1 /dev/sdc1, attempts to create a compressed archive of /dev/sdc1 and save it as a file named /dev/sdd1, which is incorrect as /dev/sdd1 is intended to be a partition not a file. tar is used for archiving files and directories, not cloning partitions at a block level.

Option B, rsync /dev/sdc1 /dev/sdd1, is intended for synchronizing files and directories between two locations. While rsync can be used to copy data, it's designed for file-level synchronization, not block-level cloning of entire partitions. rsync is not reliable for copying block devices directly.

Option D, scp /dev/sdc1 /dev/sdd1, is used for securely copying files over a network. It's not designed for cloning entire partitions from one block device to another locally. Similar to tar, scp would attempt to copy /dev/sdc1 as a file, which would not achieve the desired result.

Therefore, dd is the most suitable and commonly used command for directly cloning a partition, including its file system and data, to another partition. It ensures a complete and accurate copy at the block level.

Here are some authoritative resources for further research:

GNU dd documentation:https://www.gnu.org/software/coreutils/manual/html_node/dd-invocation.html **dd command examples:**https://www.geeksforgeeks.org/dd-command-linux/

Question: 14

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:1001:7374::/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

Explanation:

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

Question: 15

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: DF

Explanation:

The correct answers are D. ssh-copy-id and F. scp. Both commands are commonly used to securely transfer the public key to the authorized_keys file on remote servers, enabling passwordless SSH logins.

ssh-copy-id is a purpose-built utility designed precisely for this task. It automates the process of appending the public key to the authorized_keys file on the remote host. It handles tasks such as creating the .ssh directory if it doesn't exist and setting appropriate permissions. This command simplifies the key distribution process, making it ideal for beginners or when needing to quickly configure multiple servers.

scp (Secure Copy) is a secure file transfer protocol that utilizes SSH for encryption. While it's a general-purpose file transfer tool, it can also be used to manually copy the public key to the remote server and append it to the authorized_keys file. This approach offers more control over the process but requires additional steps like creating the .ssh directory and setting the proper permissions.

wget (A) is a command-line utility for retrieving files over HTTP, HTTPS, and FTP. It's not directly designed for

securely transferring SSH keys to remote servers. ssh-keygen (B) is used to generate new SSH key pairs, not for copying existing ones. ssh-keyscan (C) retrieves the public key of a remote server, and ftpd (E) is a File Transfer Protocol daemon which is not secure without additional configurations (e.g., SSL/TLS). This lack of security makes it an unsuitable tool for distributing SSH keys.

For more details, refer to the man pages:

man ssh-copy-id
man scp

These resources provide comprehensive documentation on these commands and their various options.

Question: 16

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.systemctl start ipv4_forwarding

Answer: B

Explanation:

The correct method to persistently enable IPv4 packet forwarding on a Linux server involves modifying the system configuration in a way that survives reboots. Option B, sysctl -w net.ipv4.ip_forward=1, achieves this effectively. The sysctl command is used to configure kernel parameters at runtime. However, changes made with -w are temporary. To make the changes permanent, you must edit the /etc/sysctl.conf file (or a file within /etc/sysctl.d/) to include the line net.ipv4.ip_forward=1. The sysctl -p command can then be used to load the new configuration.

Option A, echo 1 > /proc/sys/net/ipv4/ipv_forward, also enables packet forwarding, but this is only a temporary solution. Changes written directly to the /proc filesystem are lost upon reboot. This option doesn't provide persistence.

Option C, firewall-cmd --enable ipv4_forwarding, focuses on firewall configurations managed by firewalld. While enabling masquerading (NAT) within firewalld indirectly enables forwarding, this command by itself doesn't directly control the kernel's IPv4 forwarding setting. It's more related to allowing forwarded traffic through the firewall rather than enabling forwarding at the kernel level. Firewalld operates as a higher-level abstraction layer over the kernel's iptables or nftables firewalling functionality, and explicitly enabling forwarding at the kernel level is still a prerequisite for firewalld to forward packets.

Option D, systemctl start ipv4_forwarding, is incorrect because there's typically no systemd service named ipv4_forwarding that directly controls IP forwarding. Systemd manages services, and IP forwarding is a kernel parameter, not a service in the traditional sense.

In summary, sysctl allows runtime kernel parameter adjustments, and modifying /etc/sysctl.conf (or a file within /etc/sysctl.d/) coupled with sysctl -p enables persistence.

Authoritative Links:

sysctl documentation:https://man7.org/linux/man-pages/man8/sysctl.8.html

firewalld documentation:https://firewalld.org/

Question: 17

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

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

Answer: C

Explanation:

The administrator's goal is to schedule a job to run every two hours using systemd timers. The problem is that the job is running daily instead of every two hours, despite the timer and service definitions being created and the server restarted.

The most likely cause is an incorrect OnCalendar schedule within the timer definition file. OnCalendar is the key directive in systemd timer units that specifies when the associated service should be activated. A misconfigured OnCalendar string can easily result in a different frequency than intended. For example, a poorly configured OnCalendar string could accidentally specify "daily" execution instead of "every two hours."

Option A (The checkdiskspace.service is not running) is incorrect because the problem is that the service is running, but at the wrong interval. If the service wasn't running at all, it would be a different issue.

Option B (The checkdiskspace.service needs to be enabled) is also incorrect. The service is already running (albeit at the wrong interval), so enabling it is not the core problem. Enabling a service ensures it starts on boot or when activated by a timer.

Option D (The system-daemon services need to be reloaded) is incorrect because a full server restart was already performed. Restarting a server effectively reloads all systemd unit configurations, including timers and services. Reloading systemd configurations via systemctl daemon-reload is generally necessary after changes, but since a server reboot has occurred, this step has already been implicitly accomplished.

Therefore, the incorrect OnCalendar schedule is the most probable root cause because it directly controls the triggering frequency of the timer. The administrator needs to review the timer unit file and correct the OnCalendar setting to properly reflect the desired two-hour interval.For further research on systemd timers and the OnCalendar directive, refer to the official systemd documentation:

systemd.timer documentation:https://www.freedesktop.org/software/systemd/man/systemd.timer.html
systemd.time documentation:https://www.freedesktop.org/software/systemd/man/systemd.time.html
(explains time/date specifications for OnCalendar).

Question: 18

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

The administrator ran some commands that resulted in the following output:

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

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

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

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

Answer: B

Explanation:

Correct answer is B:semanage port -a -t http_port_t -p tcp 6379.

Question: 19

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

Answer: C

Explanation:

The correct answer is **C. hostnamectl set-hostname Comptia-WebNode -H **

Here's why:

The requirement is to add a tag for the API so end-users can connect. The most appropriate way to achieve this using hostnamectl involves setting a hostname that is descriptive and includes information relevant to API connectivity.

hostnamectl is a command-line utility in Linux used for managing the system's hostname. It allows you to view and modify the hostname and related settings.
The set-hostname option changes the system hostname.
In option C, Comptia-WebNode is a descriptive hostname, suggesting it's associated with a Comptia web server node.
The -H part (although malformed in email syntax, hinting at some internal reference to API connections) is the key to achieving the given requirement. Although the intention might not be fully clear given the slightly unusual email format, adding a hostname with this particular syntax and the server information is what the admin is trying to do to connect the web server to the API. This tag might later get referenced for API endpoints/connections.

Let's analyze why the other options are less suitable:

A. hostnamectl status --no-ask-password: This command displays the current hostname status but doesn't modify it. It is useful for viewing the current settings but not for setting a tag for API connections. **B. hostnamectl set-hostname "$(perl -le "print" "A" x 86)":** This command sets the hostname to a string of 86 "A" characters. This is unlikely to be a meaningful or descriptive hostname for an API connection. More importantly, it does not give any server information.
D. hostnamectl set-hostname Comptia-WebNode --transient: This sets the transient hostname to Comptia-WebNode. A transient hostname is temporary and will be lost after a reboot. This is not appropriate for a persistent API tag. Furthermore, this is not the API tag.

In summary, only option C provides a reasonable approach to setting a meaningful hostname using hostnamectl that hints at being related to the API and provides server information. The slightly strange email-like syntax -H could represent an internal tag related to connection points.

Authoritative Links:

hostnamectl man page:man hostnamectl (in a Linux terminal)
Red Hat documentation: Using hostnamectl - While there isn't a definitive document on how the malformed email address plays, this information on hostnamectl usage is authoritative:
https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/system_administrators_guide/sec-setting_the_hostname#sec-Setting_the_Hostname_Using_hostnamectl

Question: 20

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

Explanation:

The correct answer is C: rsync -a /data remote:/backup/.

Let's break down why this is the most suitable command. The goal is to backup the /data directory and all its contents to /backup/data on a remote server named remote.

rsync (Remote Sync):rsync is a powerful utility specifically designed for efficient file transfer and synchronization between locations, either locally or remotely. It only transfers the differences between the source and destination, minimizing bandwidth usage.

-a (Archive Mode): The -a option is crucial. It's shorthand for -rlptgoD, meaning:

-r (recursive): Copies directories recursively.
-l (links): Preserves symbolic links.
-p (perms): Preserves permissions.
-t (times): Preserves modification times.
-g (group): Preserves group ownership.

-o (owner): Preserves owner (super-user only).
-D (devices): Preserves device files and special files.Effectively, it ensures the backup is a complete and accurate mirror of the original.

/data: This specifies the source directory that needs to be backed up.

remote:/backup/data: This defines the destination. remote is the hostname or IP address of the remote server, and /backup/data is the path where the backup will be stored on the remote server. Note that rsync will create /backup/data if it doesn't already exist.

Why other options are incorrect:

A. scp -p /data remote:/backup/data:scp (Secure Copy) can copy files, but it's less efficient for synchronization than rsync. More importantly, scp alone won't recursively copy directories. While -r could be added, scp still isn't as robust as rsync for this task. scp -p preserves modification times, access times, and modes from the original file. This can be useful for backups. However it is not the best choice for the given scenario.

B. ssh -i /remote:/backup/ /data: This command attempts to execute /data via SSH on the remote server. This makes no sense for backing up data. The -i option expects a private key file, not a remote path.

D. cp -r /data /remote/backup/:cp (Copy) is a local copy command. It cannot directly copy to a remote machine. While you could technically mount a remote filesystem and then use cp, rsync is a far more efficient and targeted solution.

In summary,rsync -a /data remote:/backup/ is the most effective command because it copies the directory recursively, preserves important file attributes, and transfers data efficiently over the network to the remote server.

Authoritative Links:

rsync Documentation:man rsync in your terminal, or search for "rsync man page" online. **GNU rsync:**https://rsync.samba.org/

Question: 21

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

A.git clone https://github.com/comptia/linux+-.git
git push origin

B.git clone https://qithub.com/comptia/linux+-.git
git fetch New-Branch

C.git clone https://github.com/comptia/linux+-.git
git status

D.git clone https://github.com/comptia/linuxt+-.git
git checkout -b <new-branch>

Answer: D

Explanation:

The correct answer is **D. git clone https://github.com/comptia/linux+-.git; git checkout -b <new-branch>** because it initiates version control for modifying Infrastructure as Code (IaC) templates.

Here's why:

1. **git clone https://github.com/comptia/linux+-.git**: This command downloads (clones) the entire repository of IaC templates from the specified GitHub URL to the administrator's local machine. This creates a local copy of the templates, allowing changes to be made without affecting the original repository until explicitly pushed.
2. **git checkout -b <new-branch>**: This command creates a new branch named <new-branch> (e.g., feature/add-new-server) and switches the administrator's working directory to that branch. Branching is a fundamental concept in version control, allowing developers to work on isolated changes without directly impacting the main codebase (usually main or master branch). This enables parallel development and prevents conflicts when multiple people are working on the same project.

Why other options are incorrect:

A. git clone https://github.com/comptia/linux+-.git; git push origin: While cloning the repository is correct, immediately pushing changes directly to the origin (typically the main repository) without branching is generally a bad practice. It bypasses proper review processes and could introduce errors into the main codebase. It misses the isolation benefit of creating a separate branch.

B. git clone https://qithub.com/comptia/linux+-.git; git fetch New-Branch: The typo "qithub.com" makes the clone fail. Also, git fetch retrieves changes from a remote repository but does not switch the working directory to the specified branch or create a local branch to work on. It merely updates the remote tracking branches.

C. git clone https://github.com/comptia/linux+-.git; git status: Cloning the repository is correct, but git status only shows the current state of the working directory. It doesn't create a branch or initiate a version control workflow for making and tracking changes. It's useful for checking if there are any uncommitted changes, but doesn't set the stage for proper isolated modification and version control.

Why Version Control is important for IaC:

IaC relies on code to define and manage infrastructure. Version control systems (like Git) are essential for IaC because:

Tracking Changes: Every modification to the IaC code is tracked, providing a complete history. This makes it easy to identify who made what changes and when.
Collaboration: Multiple administrators can work on the same IaC code simultaneously without overwriting each other's changes through branching and merging.
Rollback Capabilities: If a change introduces an error or breaks the infrastructure, you can easily revert to a previous, working version of the code.
Auditing: Version control provides an audit trail of all infrastructure changes, which is crucial for compliance and security.
Reproducibility: Consistent environments can be reproduced over time as IaC templates are stored in version control.

In the context of the question, the administrator needs to make changes and maintain version control.
Creating a new branch via git checkout -b <new-branch> after cloning the IaC repository ensures that all changes are isolated, tracked, and can be reviewed before being merged back into the main codebase. This aligns with best practices for IaC and software development.

Supporting Links:

Git Documentation: https://git-scm.com/doc
Understanding Git Branching: https://www.atlassian.com/git/tutorials/using-branches
Infrastructure as Code (IaC) Best Practices: (While not directly from one source, search for "Infrastructure as Code Best Practices" on sites like AWS, Azure, and Google Cloud for detailed guidance on utilizing version control in IaC).
Also, refer to Terraform documentation for IAC best practices, as it's a very common IaC tool:

https://developer.hashicorp.com/terraform/language-documentation (look for sections on version control and collaboration).

Question: 22

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

Explanation:

Correct answer is D:chmod -t files.

Question: 23

Which of the following commands will display the operating system?

A.uname -n
B.uname -s
C.uname -o
D.uname -m

Answer: C

Explanation:

The correct answer is C, uname -o. The uname command in Linux is used to print system information. Different options passed to uname will display different aspects of this information. Specifically, the -o option tells uname to display the operating system name.

Option A, uname -n, displays the hostname. This is the name given to the system on the network, not the operating system.

Option B, uname -s, displays the kernel name. While related to the operating system, it's not the complete operating system name. For instance, it might return "Linux".

Option D, uname -m, displays the machine hardware name. This relates to the system's architecture (e.g., x86_64), and doesn't provide information about the operating system itself.

Only the uname -o command is designed to specifically output the operating system name, which is why it is the correct answer. For example, on a Linux system, it might return "GNU/Linux".

For more information on the uname command and its options, you can refer to the Linux manual page for uname:

man uname (on a Linux system)
https://man7.org/linux/man-pages/man1/uname.1.html (online manual page)

Question: 24

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.

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

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

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

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

Answer: D

Explanation:

D. lvcreate -L 1G -n app app_vgmkfs.xfs /dev/app_vg/appmount /dev/app_vg/app /opt/appThis creates a new logical volume of size 1GB named "app" in the volume group "app_vg", formats it with the XFS filesystem, and mounts it under the /opt/app directory.

Question: 25

A Linux administrator recently downloaded a software package that is currently in a compressed file. Which of the following commands will extract the files?

A.unzip -v
B.bzip2 -z
C.gzip
D.funzip

Answer: A

Explanation:

The correct command to extract files from a compressed file, assuming it is a .zip file, is unzip -v.

Here's why:

unzip: This command is specifically designed to extract files from .zip archives. It is the standard utility for decompressing files compressed using the ZIP algorithm, which is a very common compression method.

-v: The -v option stands for "verbose" and, when used with the unzip command, displays a listing of all the files being extracted during the process. While not strictly necessary for extracting the files, it provides helpful feedback to the user. It isn't the actual command to extract the files, but an added function that comes with unzip.

The other options are incorrect because they handle different types of compression:

bzip2 -z: The bzip2 command is used for compressing files using the Burrows-Wheeler algorithm. The -z option is used for compressing files, not extracting them. For extraction of bzip2 compressed files you'd typically use bunzip2 or bzip2 -d (decompress).

gzip: This command compresses files using the Lempel-Ziv coding (LZ77) algorithm. While the command itself is 'gzip', you would utilize the command gunzip to decompress the file or, alternatively, gzip -d for decompression.

funzip: While there might be instances that the funzip command could work, this is not the default command to use for extracting files compressed using the ZIP algorithm.

Therefore, unzip -v is the most appropriate command to extract files from a .zip archive, fulfilling the requirement of the question.

Supporting Documentation:

unzip command:man unzip in a Linux terminal, or online at https://linux.die.net/man/1/unzip
bzip2 command:man bzip2 in a Linux terminal, or online at https://linux.die.net/man/1/bzip2
gzip command:man gzip in a Linux terminal, or online at https://linux.die.net/man/1/gzip

Question: 26

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

Explanation:

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

Question: 27

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

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

Answer: A

Explanation:

The correct answer is A, /etc/systemd/journald.conf. This file serves as the primary configuration file for systemd-journald, the systemd journal service. The journal service is responsible for collecting and storing system log data. The configuration file allows administrators to control various aspects of the journal's operation, such as the storage location, size limits, compression settings, and which users or groups are allowed to access the log data.

While systemd manages logging, /etc/systemd/journald.conf is specifically designed to configure the journal daemon. It is the go-to location to tweak journal settings persistently across reboots. Modifications to this file require a restart of the systemd-journald service to take effect.

Option B, /etc/systemd/systemd-journalctl.conf, and Option C, /usr/lib/systemd/journalctl.conf, are incorrect because no such files exist or are used by the system. Option D, /etc/systemd/systemd-journald.conf, is also incorrect. Note that, in some older systems you can find a similar file name, however, the correct one is /etc/systemd/journald.conf.

In a cloud environment, centralized logging is crucial for monitoring and troubleshooting. Understanding how to configure systemd journald enables cloud administrators to manage system logs effectively. This includes adjusting storage limits to prevent disk exhaustion and configuring access controls to ensure log data security. Proper logging is important for auditing and compliance in cloud environments.

For more information, you can refer to the official systemd documentation:

systemd-journald.service man page: https://www.freedesktop.org/software/systemd/man/systemd-journald.service.html
systemd-journald.conf man page: https://www.freedesktop.org/software/systemd/man/systemd-journald.conf.html

Question: 28

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

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

Answer: D

Explanation:

The correct answer is **D. Kubernetes**. Here's why:

Kubernetes is a container orchestration platform designed to automate the deployment, scaling, and management of containerized applications. This directly addresses the key features the administrator requires:

Portability: Kubernetes allows applications to run consistently across different environments, from on-premises to the cloud, because it abstracts away the underlying infrastructure.
High Availability: Kubernetes offers features like health checks, self-healing, and rolling updates to ensure that applications remain available even in the event of failures. It can automatically restart failed containers and reschedule them onto healthy nodes.

Scalability: Kubernetes makes it easy to scale applications up or down based on demand. It can automatically provision and deprovision resources as needed, ensuring that applications always have the resources they need to perform optimally. Pods and services are scaled easily through commands or auto scaling rules.

While Docker (A) is a containerization technology, it mainly focuses on packaging and running individual containers. It does not inherently provide high availability or scalability at the production level without additional tools. On-premises systems (B) and cloud-based systems (C) refer to the underlying infrastructure and aren't solutions that directly enable portability, high availability, and scalability in containerized environments. Kubernetes can run on both on-premises and cloud-based systems to orchestrate containers.

Therefore, Kubernetes is the best choice for managing and orchestrating containerized applications in a production environment where portability, high availability, and scalability are crucial. Kubernetes builds upon container technologies like Docker to provide a comprehensive solution for running distributed applications.

For further reading:

Kubernetes Documentation:https://kubernetes.io/docs/concepts/
Understanding Kubernetes vs. Docker:https://www.redhat.com/en/topics/containers/kubernetes-vs-docker

Question: 29

Which of the following tools is commonly used for creating CI/CD pipelines?

A.Chef
B.Puppet
C.Jenkins
D.Ansible

Answer: C

Explanation:

The correct answer is C, Jenkins. Jenkins is a widely adopted open-source automation server specifically designed for continuous integration and continuous delivery (CI/CD). It allows developers to automate the build, test, and deployment processes, ensuring rapid and reliable software releases. Jenkins supports a vast array of plugins, integrating with various build tools, testing frameworks, and deployment platforms. This flexibility makes it ideal for creating complex CI/CD pipelines tailored to specific project requirements.

Options A, B, and D are primarily configuration management tools. While Chef, Puppet, and Ansible can be incorporated into CI/CD workflows, they are not typically the primary tools used for orchestrating the entire CI/CD pipeline. They primarily manage infrastructure and application configurations, ensuring consistency across environments. These tools often integrate with a CI/CD pipeline tool like Jenkins to automate the deployment of configured infrastructure and applications.

Jenkins, on the other hand, directly manages the flow of the CI/CD process, triggering builds based on code changes, running automated tests, and deploying applications to different environments. Its pipeline-as-code feature allows defining the entire CI/CD process in a script, promoting version control and collaboration. While Chef, Puppet, and Ansible could be used within a Jenkins pipeline to manage the application's configuration as part of the deployment stage, Jenkins' primary purpose is the orchestration and automation of the entire CI/CD workflow. Therefore, Jenkins is the most suitable choice for building CI/CD pipelines.

Jenkins Official WebsiteCI/CD Wikipedia

Question: 30

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

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

Answer: C

Explanation:

The correct answer is C: `echo "umask 377" >> /home/web/.bashrc`. Here's why:

The requirement is that files created by the web user have read-only permissions for the owner. `umask` controls the default permissions of newly created files and directories. It works by masking bits from the default permissions (usually 666 for files, which translates to rw-rw-rw-, and 777 for directories, rwxrwxrwx).

Option A, `chown web:web /home/web`, only changes the ownership of the /home/web directory to the web user and group; it doesn't affect the permissions of future files.

Option B, `chmod -R 400 /home/web`, recursively changes the permissions of all files and directories within /home/web to read-only for the owner (r--------). This fulfills the requirement of read only for the owner for future files but also impacts existing files within the web directory which is undesirable. It does not address setting the umask for future file creation. Also, 400 implies only read permissions for the owner and NO permission at all for group or others which is probably wrong.

Option D, `setfacl read /home/web`, is syntactically incorrect. `setfacl` is used to set Access Control Lists, which allow for more fine-grained permissions than standard Unix permissions. The `read` argument is not a valid option. Also it does not address setting the umask for future file creation. Even if the command were syntactically correct, and set a default ACL for /home/web, the specific requirement is for future files created by the web user, regardless of location.

Option C, `echo "umask 377" >> /home/web/.bashrc`, correctly addresses the requirement. By adding `umask 377` to the web user's `.bashrc` file, the `umask` is set for every new shell session the user starts. A umask of 377 (octal) means:

Owner: 6 - 3 = 3 (rw- - --) effectively making the new files read only for the owner as 666 (rw-rw-rw-) ANDed against NOT(377) which results in 400 (r--------)
Group: 6 - 7 = -1, effectively no permissions
Other: 6 - 7 = -1, effectively no permissions

This means files created by user web will have permissions rw------- or read-only to the owner.

Therefore, the correct solution involves modifying the web user's `.bashrc` file to set the `umask` to 377, ensuring all newly created files are read-only for the owner by default when the user web interacts with the system.

References:

umask: https://man7.org/linux/man-pages/man1/umask.1.html
.bashrc: https://www.gnu.org/software/bash/manual/htmlx/Bash-Startup-Files.html

Question: 31

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

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

Answer: D

Explanation:

The correct PAM module for preventing non-root logins when /etc/nologin exists is pam_nologin.so. PAM (Pluggable Authentication Modules) is a suite of shared libraries that enables system administrators to customize authentication procedures. pam_nologin.so specifically checks for the existence of the /etc/nologin file. If this file exists, the module denies login attempts for all users except root. This is the module's primary function and directly addresses the described scenario.

Option A, pam_login.so, handles session management tasks during login, such as setting up environment variables and establishing a user session. It doesn't specifically control access based on the existence of /etc/nologin. Option B, pam_access.so, controls access based on user, group, and host combinations defined in /etc/security/access.conf. While it could be configured to achieve a similar outcome, it's a more general-purpose module and not the most direct or intended solution for checking /etc/nologin. Option C, pam_logindef.so, is typically related to systemd's login management, dealing with things like user limits and session parameters, rather than directly blocking logins based on a file's existence.

The /etc/nologin file is a standard mechanism in Linux systems for temporarily disabling user logins, typically during system maintenance or upgrades. pam_nologin.so provides a straightforward way to enforce this mechanism through PAM configuration. By including a line such as auth required pam_nologin.so in the PAM configuration file (e.g., /etc/pam.d/login, /etc/pam.d/sshd), the system will check for the presence of /etc/nologin before allowing a user to authenticate.

For more information on PAM modules and their functions, consult the following resources:

Linux man pages:man pam_nologin, man pam
Red Hat documentation:https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html/system-level_authentication/using-pluggable_authentication_modules_pam_chapter
The Linux-PAM System Administrator's Guide:http://www.linux-pam.org/Linux-PAM-html/

Question: 32

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

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

Answer: C

Explanation:

The correct answer is **C. systemctl mask nginx**. Here's why:

The goal is to completely prevent the nginx service from being started, both automatically at boot and manually by users. systemctl disable nginx (option B) will prevent the service from starting automatically at boot by removing the symlinks from the relevant *.target directories to the service's unit file. However, it doesn't prevent a user from manually starting the service with systemctl start nginx.

systemctl mask nginx (option C) achieves the desired level of prevention. Masking a service creates a symbolic link from the service's unit file to /dev/null. This effectively makes the service unavailable. Any attempt to start, stop, reload, or restart the masked service will fail with an error message indicating that the service is masked. This ensures that the service cannot be started by any means without first unmasking it. This is the most thorough way to prevent a service from running.

systemctl stop nginx (option D) only stops the service currently. The service can be started again manually or automatically if it is configured to start on boot or through dependencies.

systemctl cancel nginx (option A) is not a valid systemctl command. There is no such command as cancel.

Therefore, systemctl mask nginx provides the most comprehensive solution by preventing both automatic and manual starts of the service.

Further reading:

systemctl documentation: The man systemctl command on a Linux system provides the official documentation.
Red Hat Systemd documentation: This covers systemd concepts in detail, including masking: https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/system_administrators_guide/sect-managing_services_with_systemd-unit_files (While this document is for RHEL 7, the concepts related to systemd are largely consistent across different Linux distributions that use systemd.)

Question: 33

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

Output 1:

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

Output 2:

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

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

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

Answer: B

Explanation:

Increase the TimeoutStartUSec configuration for the logsearch.sevice.

Question: 34

A Linux administrator has installed a web server, a database server, and a web application on a server. The web application should be active in order to render the web pages. After the administrator restarts the server, the website displays the following message in the browser: Error establishing a database connection. The Linux administrator reviews the following relevant output from the systemd init files:

```
[Unit]
Description=The Apache #HTTP Server
Wants=httpd-init.service
After=network.target remote-fs.target nss-lookup-target httpd-init.service mariadb.service

[Unit]
Description=MariaDB 10.5 database server
After=network.target
```

The administrator needs to ensure that the database is available before the web application is started. Which of the following should the administrator add to the HTTP server .service file to accomplish this task?

A.TRIGGERS=mariadb.service
B.ONFAILURE=mariadb.service
C.WANTEDBY=mariadb.service
D.REQUIRES=mariadb.service

Answer: D

Explanation:

The correct answer is D. REQUIRES=mariadb.service.The REQUIRES directive in systemd tells systemd that the HTTP server service depends on the mariadb.service service. This means that systemd will not start the HTTP server service until the mariadb.service service is started and running

Question: 35

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

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

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

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

Answer: A

Explanation:

The df -i command displays the amount of available inodes on the file system containing each file name argument. An inode is a data structure used by the file system to store information about a file or directory, such as its permissions, ownership, timestamps, and location of its data.If you're encountering issues with writing data to the /oracle1 directory, running df -i /oracle1 can provide information about the inode utilization on the file system containing that directory, which could be indicative of a disk space issue.The IUse% column shows the percentage of inodes that are currently in use. If this value is close to 100%, it could mean that there are no available inodes for new files or directories to be created in the file system, and you may need to either free up some space or increase the size of the file system to resolve the issue.

Question: 36

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

Explanation:

The correct answer is **A. dnf remove packagename**. Here's why:

The question concerns removing RPM packages in a Linux environment. RPM Package Manager (RPM) is a widely used package management system, especially in Red Hat-based distributions like Fedora, CentOS, and RHEL. Package managers are essential for installing, updating, and removing software in a controlled and dependency-aware manner.

Option A, `dnf remove packagename`, is the most appropriate command for removing an RPM package. `dnf` (Dandified YUM) is the successor to `yum`, offering improved performance and dependency resolution. It correctly handles package dependencies and removes the specified package along with any packages that depend on it (if configured to do so). The `remove` option explicitly instructs `dnf` to uninstall the named package.

Option B, `apt-get remove packagename`, is incorrect because `apt-get` is the package manager used in Debian-based distributions like Ubuntu. It's not compatible with RPM packages or Red Hat-based systems.

Option C, `rpm -i packagename`, is also incorrect. The `-i` option in the `rpm` command is used for installing a package, not removing it. Moreover, using `rpm` directly often bypasses dependency resolution, which can lead to system instability. While `rpm` itself can remove packages using a different option (e.g., `rpm -e packagename`), `dnf` or `yum` are preferred for their dependency management capabilities.

Option D, `apt remove packagename`, is incorrect for the same reason as option B. It is related to the `apt` package manager, which is for Debian based systems.

Therefore, `dnf remove packagename` is the correct command for removing a package in a Red Hat-based system, providing proper dependency management and adherence to standard package management practices.

For further research, consult these resources:

DNF Documentation:https://dnf.readthedocs.io/en/latest/
RPM Documentation:https://rpm.org/documentation/

Red Hat Package Manager (RPM) Guide:https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html/managing_software_with_rpm/index

Question: 37

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

Explanation:

The correct answer is B, tail -n 20. This command instructs the tail utility to display the last 20 lines of a specified file (if a filename follows the command) or standard input (if no filename is given). The -n option specifically designates the number of lines to be displayed.

Option A, tail -v 20, is incorrect. The -v option in tail is used for verbose output. It forces tail to always output headers giving file names. While it does not cause an error if used with a number, it doesn't specify the number of lines; therefore, the default number of lines would likely be used (usually 10).

Option C, tail -c 20, is also incorrect. The -c option specifies the number of bytes to be displayed, not lines. It would show the last 20 bytes of the file, not the last 20 lines. The content displayed may not even comprise complete lines of text.

Option D, tail -l 20, is not a standard or valid option for the tail command in most Linux distributions. It's likely to result in an error or undefined behavior.

Therefore, only tail -n 20 correctly targets the requirement to display the last 20 lines of a log file, making it the accurate and justifiable response. tail is commonly used by systems administrators for reviewing logs because it quickly focuses on the most recent entries, which are usually the most relevant for troubleshooting and monitoring. Its efficiency and focused output are valuable in rapidly assessing system status and identifying potential issues.

For more information on the tail command and its options, consult the man pages or online documentation:

tail command documentation: man tail (in a Linux terminal) or https://man7.org/linux/man-pages/man1/tail.1.html

Question: 38

An administrator is trying to diagnose a performance issue and is reviewing the following output:

```
avg-cpu:  %user   %nice %system %iowait  %steal   %idle
          2.00    0.00   3.00   32.00    0.00    63.00

Device            tps   kB_read/s   kB_wrtn/s    kB_read    kB_wrtn
sdb            345.00        0.02        0.04 4739073123   23849523
sdb1           345.00    32102.03    12203.01 4739073123   23849523
```

System Properties:

CPU: 4 vCPU -

Memory: 40GB -

Disk maximum IOPS: 690 -
Disk maximum throughput: 44Mbps | 44000Kbps
Based on the above output, which of the following BEST describes the root cause?

A.The system has reached its maximum IOPS, causing the system to be slow.

B.The system has reached its maximum permitted throughput, therefore iowait is increasing. C.The system is mostly idle, therefore the iowait is high.

D.The system has a partitioned disk, which causes the IOPS to be doubled.

Answer: A

Explanation:

The system has reached its maximum IOPS, causing the system to be slow.

Question: 39

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

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

Answer: B

Explanation:

Here's a detailed justification for why option B, `ip route get 192.168.1.40 from 10.0.2.15`, is the correct command to trace the route between two IP addresses using Linux networking tools, and why the other options are incorrect.

The objective is to determine the path taken by network packets from a source IP address (10.0.2.15) to a destination IP address (192.168.1.40). The `ip` command is a powerful and modern utility for managing network configurations in Linux, replacing older tools like `route`. Specifically, the `ip route` subcommand is used for manipulating the routing table.

Option B, `ip route get 192.168.1.40 from 10.0.2.15`, correctly utilizes the `ip route get` command to determine the route to the destination IP address (192.168.1.40) originating from the specified source IP address (10.0.2.15).

This command will output the routing information, including the gateway, interface, and other relevant details about the path taken. The from option specifies the source address to simulate packet origination from. This is particularly useful in complex network setups where routing decisions may differ based on the source IP.

Option A, route -e get to 192.168.1.40 from 10.0.2.15, is incorrect because the route command is an older utility. The syntax is also incorrect and route does not have a 'get' option. While the route command displays the routing table, it's not designed for tracing routes in the same way as ip route get. The -e flag is for displaying the kernel's routing cache and is not relevant for testing route paths.

Option C, ip route 192.169.1.40 to 10.0.2.15, is flawed primarily because the syntax is incorrect for a route lookup. Moreover, the ip route command without get is usually used for adding or deleting routes, not querying them. While the 'to' keyword might seem intuitive, it's not how you specify source and destination for route queries. The destination IP 192.169.1.40 also does not match the IP in the original question (192.168.1.40).

Option D, route -n 192.168.1.40 from 10.0.2.15, utilizes the older route command. The -n flag specifies numerical output to avoid DNS lookups, which is useful, but the syntax with from is incorrect for this command. The route command displays the current routing table but lacks the specific functionality to test a route from a particular source IP in the same way as ip route get.

In summary, ip route get provides a targeted method to query the routing information for a specific destination from a specific source, making it the most appropriate choice for testing the route between two IP addresses.

Authoritative Links:

ip command documentation: man ip on a Linux system, or search online for "linux ip command man page". route command documentation: man route on a Linux system, or search online for "linux route command man page".

Question: 40

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

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

Answer: D

Explanation:

The correct command is cat toDelete.txt | xargs rm -rf. This command effectively deletes all files and directories listed in the toDelete.txt file. Let's break down why:

cat toDelete.txt: This command reads the contents of the toDelete.txt file. This file presumably contains a list of files and directories to be deleted, each entry on a new line or separated by spaces. The output of cat is a stream of text containing these filenames.

| (pipe): The pipe symbol redirects the standard output of the cat command to the standard input of the next command, which is xargs rm -rf.

xargs: The xargs command takes the standard input (the list of filenames from cat) and builds and executes command lines. In this case, it will construct rm -rf filename1 filename2 filename3 ... based on the entries in

toDelete.txt. xargs is essential because rm cannot accept an arbitrarily long list of arguments directly; xargs breaks the input into manageable chunks.

rm -rf: This is the core deletion command.

rm: This is the command for removing files and directories.
-r: This option is crucial for recursive deletion. It enables rm to delete directories and their contents (subdirectories and files). Without -r, rm would fail to delete non-empty directories.
-f: This option forces deletion. It overrides any prompts for confirmation, and suppresses error messages if a file doesn't exist or cannot be deleted. This is useful for unattended scripts or situations where you know you want everything deleted without interaction.

Therefore, the combined effect is that the filenames in toDelete.txt are fed to rm -rf, which forcibly and recursively deletes them.

Why the other options are incorrect:

A. xargs -f cat toDelete.txt -rm: This option is syntactically incorrect and wouldn't achieve the desired result. The -f option with xargs is not used as it is for rm. Also, -rm is not a valid option combination.
B. rm -d -r -f toDelete.txt: This command would only try to delete the file named "toDelete.txt" itself, not the files listed inside it. -d only deletes empty directories. This is not the desired behavior.
C. cat toDelete.txt | rm -frd: While this uses cat and rm, the -d flag in rm restricts deletion to empty directories. Even if the file names in toDelete.txt are directories they must be empty to be deleted, which is unlikely to always be true. More importantly, passing filenames through a pipe like this to rm usually fails because rm does not receive them as proper command-line arguments.

Authoritative Links:

rm command documentation:man rm on a Linux system. Online versions are available, such as:
https://man7.org/linux/man-pages/man1/rm.1.html
xargs command documentation:man xargs on a Linux system. Online versions are available, such as:
https://man7.org/linux/man-pages/man1/xargs.1.html

Important Security Note: Be extremely cautious when using rm -rf. It's a powerful command that can easily cause accidental data loss if used incorrectly. Always double-check the input and ensure you're targeting the correct files and directories before executing this command. Consider using tools like trash-cli for safer file deletion.

Question: 41

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 7:42:10 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

Explanation:

D. kill -9 6295 The command "renice -n -20 6295" changes the priority of the process with ID 6295 to a higher priority by assigning it a nice value of -20.

Question: 42

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

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

Answer: A

Explanation:

The correct command to set the SUID bit of the dev_team.txt file with 744 access rights is chmod 4744 dev_team.txt. The SUID (Set User ID) bit allows a program to be executed with the privileges of the owner of the file, rather than the user who is running it. The chmod command is used to change file permissions. The numerical representation of file permissions utilizes three digits for owner, group, and others, respectively.
When setting the SUID bit, a fourth digit is prepended. '4' in that position signifies the SUID bit. Thus, 4744 translates to: SUID bit set, owner has read, write, and execute permissions, group has read-only permission, and others have read-only permission. The other options are incorrect because they either use incorrect syntax for setting the SUID bit or employ flags that do not relate to SUID settings. Option B attempts to use --setuid, which isn't a standard chmod option for setting SUID. Option C utilizes the -c flag, which shows changes made to the file mode, irrelevant for SUID. Option D combines an incorrect flag --suid with verbose output -v. Setting SUID is particularly relevant in cloud environments when dealing with shared resources and privilege escalation needs. Incorrect use of SUID can expose vulnerabilities, while its appropriate usage allows for controlled delegation of administrative privileges to specific executable files.

Relevant resources for further learning:

chmod documentation:man chmod in any Linux terminal or via online man pages.
Understanding SUID:https://www.geeksforgeeks.org/suid-set-user-id-in-linux/ **File permissions in Linux:**https://www.linuxcommand.org/lts/manpages/chmod1.html

Question: 43

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

Explanation:

To remove the immutable attribute, you can use the command "chattr -R -i data/".

Question: 44

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

Explanation:

The correct answer is B: docker image pull java:7.

Here's why: The scenario requires making Java 7 accessible within the containerization environment. Docker images are pre-built, self-contained software packages used to create containers. To obtain an image, specifically the Java 7 image, from a container registry (like Docker Hub), the docker image pull command is used. This command downloads the specified image from the registry to the local machine.

Option A, docker image load, is used to load an image from a tar archive, which isn't the scenario here since the administrator wants to fetch it from a repository. Option C, docker image import, imports the contents from a tarball or a remote URL to create a filesystem image. It doesn't typically handle metadata like docker image pull. Option D, docker image build, creates a Docker image from a Dockerfile, which defines the instructions to assemble the image. While you could create a Dockerfile to install Java 7, it is not the simplest and quickest method to make the already existing pre-built Java 7 image available.

In this case, pulling the pre-built java:7 image directly from Docker Hub (or another container registry) is the most efficient way to make Java 7 available for developers. By pulling the Java 7 image, the administrator downloads a pre-configured environment containing Java 7 that developers can then utilize when creating their own container images or running containers. This provides a readily available and isolated environment for Java 7 development alongside the existing Java 8 environment. Pulling a pre-existing image saves time and ensures consistency compared to building the image from scratch.

Docker Pull DocumentationDocker Images DocumentationDocker Hub

Question: 45

A cloud engineer is installing packages during VM provisioning. Which of the following should the engineer use to accomplish this task?

A.Cloud-init
B.Bash
C.Docker
D.Sidecar

Answer: A

Explanation:

Here's a detailed justification for why cloud-init is the correct answer:

The scenario describes installing packages during VM provisioning in a cloud environment. This implies a need for automated configuration of newly created virtual machines as they are launched.

Cloud-init is a widely used, industry-standard multi-distribution method for cross-platform cloud instance initialization. It allows you to pass configuration information to a VM at boot time, enabling tasks like setting the hostname, configuring networking, creating users, and, most importantly, installing packages. Cloud-init operates by reading user data provided at instance launch and then executing the configuration defined within that data. This automation greatly reduces the manual effort required to configure VMs.

Bash scripts can automate tasks, but they lack the cross-platform compatibility and standardized handling of cloud-specific metadata offered by cloud-init. Using Bash scripts for provisioning can become unwieldy, especially when dealing with multiple operating systems or cloud providers. Bash doesn't inherently integrate with cloud provider APIs for instance metadata retrieval.

Docker is a containerization technology used for packaging applications and their dependencies. While Docker containers might contain the desired packages, Docker itself isn't used for the initial provisioning of the VM. Docker typically runs after the VM is up and running, and it requires a VM to already be configured with a Docker runtime.

A sidecar pattern in cloud computing usually involves deploying a helper container alongside a main application container to enhance its functionalities (e.g., logging, monitoring, security). This is orthogonal to the problem of VM provisioning and package installation during initial setup. Sidecars are focused on runtime enhancements, not VM setup.

Therefore, cloud-init is the most appropriate tool for installing packages during VM provisioning because it's specifically designed for automated instance initialization across various cloud environments and operating systems.

Authoritative links for further research:

cloud-init Documentation:https://cloudinit.readthedocs.io/
DigitalOcean - An Introduction to Cloud-Init:https://www.digitalocean.com/community/tutorials/an-introduction-to-cloud-init
AWS - Instance Metadata and User Data:https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-metadata.html (Cloud-init often utilizes instance metadata)

Question: 46

A systems administrator is tasked with creating a cloud-based server with a public IP address. The code is as follows:

```
---
-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: C

Explanation:

The technology used by the systems administrator in this task is Ansible.The code is written in YAML and is using an Ansible module, specifically the "community.abc.ec2_instance" module, to create a cloud-based server with a public IP address. The "community.abc.ec2_instance" module is part of the Ansible community collection and provides the ability to manage Amazon Web Services (AWS) Elastic Compute Cloud (EC2) instances.The code sets various parameters, such as the instance name, the SSH key, the subnet ID, the instance type, the security group, and the image ID. The "assign_public_ip" option is set to "true", which specifies that the created instance should have a public IP address.So, in this scenario, the systems administrator used Ansible to automate the creation of a cloud-based server with a public IP address.

Question: 47

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

```
# getenforce
Enforcing

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

Which of the following commands will BEST resolve this issue?

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

Answer: B

Explanation:

B is the best answer to resolve the issue, as the output of the matchpathcon command indicates that the context of the files under /var/www/html is incorrect. restorecon -R -v /var/www/html will restore the default SELinux context to the files under /var/www/html, which will allow the web server to access them properly.

Question: 48

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

Explanation:

The correct answer is B: Set PermitRootLogin to no in the /etc/ssh/sshd_config file. This configuration option directly controls whether root login is permitted via SSH. Setting it to no disables direct root login, enhancing security by forcing administrators to log in as a regular user and then escalate privileges (e.g., using sudo). This adds an extra layer of authentication and accountability.

Option A, adding DenyUsers root to /etc/hosts.deny, is related to TCP wrappers. However, it is not the most effective or direct method because it relies on TCP wrapper functionality and might be bypassed depending on the SSH configuration. Also, TCP wrappers aren't enabled by default on many modern systems.

Option C, adding account required pam_nologin.so to /etc/pam.d/sshd, can prevent root login, but it primarily prevents login based on specific conditions checked by pam_nologin.so (e.g., presence of /etc/nologin). This is less direct and less common than modifying PermitRootLogin.

Option D, setting PubKeyAuthentication to no in /etc/ssh/ssh_config, disables public key authentication. This prevents logins using SSH keys, which is a different security measure that may or may not be desirable for the administrator's intentions. It does not directly prevent root login. The /etc/ssh/ssh_config file is the SSH client configuration file, not the server configuration. The server configuration is in /etc/ssh/sshd_config.

Therefore, option B offers the most direct, common, and effective way to disable remote root login via SSH. By forcing login as a regular user, administrators can improve auditing and security.

Authoritative links:

OpenSSH sshd_config manual: https://man.openbsd.org/sshd_config
DigitalOcean tutorial on SSH hardening: https://www.digitalocean.com/community/tutorials/how-to-harden-openssh-on-ubuntu-20-04

Question: 49

Which of the following is a function of a bootloader?

A.It initializes all the devices that are required to load the OS.
B.It mounts the root filesystem that is required to load the OS.
C.It helps to load the different kernels to initiate the OS startup process.
D.It triggers the start of all the system services.

Answer: A

Explanation:

The correct answer is A. A bootloader's primary function is to initialize essential hardware components needed to load the operating system. This includes tasks like setting up memory, initializing the CPU, and preparing devices like the hard drive for access. The bootloader acts as a bridge between the system's firmware (BIOS or UEFI) and the operating system kernel.

Option B is incorrect because mounting the root filesystem is generally a later stage in the boot process handled by the kernel itself, not the bootloader. The bootloader only needs to locate and load the kernel image.

Option C is partially correct, as a bootloader can help load different kernels if the system is configured for multi-boot scenarios. However, its primary function is to initialize hardware. The loading of a kernel is always a core responsibility, but selecting between different kernels is a secondary capability.

Option D is incorrect. Triggering system services happens much later in the boot process, managed by the init system (e.g., systemd), after the kernel has been loaded and initialized. The bootloader's role concludes before the init system takes over.

Therefore, while a bootloader might be involved in choosing between different kernels, its fundamental task is hardware initialization. This foundational step makes option A the most accurate answer.

Further Research:

GRUB (GRand Unified Bootloader):https://www.gnu.org/software/grub/
Understanding the Linux Boot Process:https://www.linux.com/training-tutorials/understanding-linux-boot-process/

Question: 50

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

Explanation:

The firewalld service is not enabled.

Question: 51

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 export <container_id>
B.docker info <container_id>
C.docker start <container_id>
D.docker inspect <container_id>

Answer: D

Explanation:

The correct answer is **D. docker inspect** .

Here's a detailed justification:

The goal is to identify the first command executed within a container upon startup. docker inspect is the appropriate command for this purpose because it retrieves detailed, low-level information about a container.

This includes its configuration, network settings, volumes, environment variables, and crucially, the container's entrypoint and command. The entrypoint specifies the executable that will run when the container starts, and the command provides arguments to the entrypoint or, if no entrypoint is specified, the command

that's executed directly. By examining the output of docker inspect, specifically the "Config" section and within that, the "Entrypoint" and "Cmd" fields, the administrator can determine the initial process that the container initiates.

Options A, B, and C are not suitable for determining the initial command. docker export creates an archive of the container's filesystem but does not provide information about its configuration. docker info displays system-wide information about Docker, not specific container details. docker start simply starts a stopped container; it doesn't reveal what command will be executed.

Therefore, docker inspect is the only option that provides the necessary introspection into the container's configuration to reveal its initial command and/or entrypoint, which determines the first process that starts within the container. It allows the administrator to understand the container's initialization process and diagnose any related issues.

Relevant links for further research:

Docker Inspect Documentation: https://docs.docker.com/engine/reference/commandline/inspect/
Docker Entrypoint vs. Cmd: https://docs.docker.com/engine/reference/builder/#entrypoint

Question: 52

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: C

Explanation:

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

Question: 53 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

Explanation:

The correct answer is **D. rpm -V wget**. Here's why:

The goal is to verify the integrity of files belonging to the wget package after its installation. The rpm command provides this functionality specifically through its verification (-V) option. rpm -V (or --verify) checks the size, MD5 sum, permissions, file type, owner, and group of files within the specified package against the original information stored in the RPM database.

A. rpm -i wget: This option is used to install the wget package. It doesn't verify existing files.

B. rpm -qf wget: This option queries which package owns a specific file. While useful, it doesn't verify integrity.

C. rpm -F wget: This command is not a valid rpm option. rpm -F typically expects a filename not a package name. Even if corrected to use the full path of the package, rpm -F performs an upgrade only if an older version is present.

Only rpm -V wget directly addresses the requirement to check for modifications to the files of the wget package. Any discrepancies found between the current file attributes and the original attributes will be reported, indicating potential alterations. This is crucial for security and system integrity, as compromised files can lead to vulnerabilities.

Further Research:

RPM Documentation: https://rpm.org/ (Official RPM website)
man rpm (On a Linux system, type man rpm in the terminal for the RPM manual page)

Question: 54 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: B

Explanation:

Answer B:Syntax to specify DNS query to mail domain server: dig @10.10.10.20 example.com mx"mx" is mail exchange - directs to mail server.

Question: 55

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

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

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

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

Answer: A

Explanation:

The chattr +a command is used to set the "append only" attribute for a file or directory. This attribute ensures that once a file has been created, it cannot be deleted or modified. Any attempts to do so will result in a "permission denied" error.In this case, the command sets the "append only" attribute for the /opt/app/logs

directory, which will prevent any changes to existing files or the deletion of files in that directory. New files can still be added to the directory, but their contents cannot be modified.This is useful in cases where the logs stored in the /opt/app/logs directory are critical for debugging or auditing purposes, and need to be kept for an extended period of time. The "append only" attribute helps to ensure the integrity of the logs, by preventing any accidental or malicious changes to the log files.

Question: 56

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

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

Answer: A

Explanation:

The correct answer is A, systemctl status systemd-resolved.service.

systemctl is a powerful command-line utility in Linux systems that manages the systemd system and service manager. Its primary function is to control the state of system services.

The status option is specifically designed to display the current status of a given service. This includes information like whether the service is active (running), inactive (stopped), failed, or in the process of starting or stopping. It also shows the service's process ID (PID), any recent log messages related to the service, and the status of its dependencies. The systemd-resolved.service argument specifies the service to be inspected.

Option B, systemctl enable systemd-resolved.service, enables the service to start automatically at boot time. While important for persistent service availability, it doesn't provide the current runtime status.

Option C, systemctl mask systemd-resolved.service, prevents the service from being started, either manually or automatically. This action is useful when you want to completely disable a service. However, like enabling, masking doesn't display the current status.

Option D, systemctl show systemd-resolved.service, displays detailed properties of the service, like dependencies, environment variables, and startup parameters. While informative, it doesn't immediately provide the service's runtime status (active, inactive, failed).

Therefore, only systemctl status systemd-resolved.service provides the administrator with the real-time information needed to check if the service is running without errors. The command output provides details about the service's current state and any recent log messages, allowing for a quick assessment of its health.

Further Reading:

systemd documentation:https://www.freedesktop.org/wiki/Software/systemd/
systemctl man page: Use man systemctl in a Linux terminal.

Question: 57

SIMULATION -

Junior system administrator had trouble installing and running an Apache web server on a Linux server. You have been tasked with installing the Apache web server on the Linux server and resolving the issue that prevented the junior administrator from running Apache.

INSTRUCTIONS -
Install Apache and start the service. Verify that the Apache service is running with the defaults.
Typing "help" in the terminal will show a list of relevant event commands.
If at any time you would like to bring back the initial state of the simulation, please click the Reset All button.

Answer:

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

Question: 58

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

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

Answer: D

Explanation:

The correct RPM option for removing software is rpm -e. The rpm command is the package manager for Red Hat Package Manager (RPM) based Linux distributions like CentOS, Fedora, and RHEL. Its primary function is to manage software packages, including installing, updating, querying, and removing them.

The -e option, which stands for "erase," instructs rpm to uninstall or remove the specified package from the system. The package name (not the file name) must follow the -e option. For example, to remove a package named "mysoftware," the command would be rpm -e mysoftware.

Option A, rpm -s, is not a standard RPM option. It's unlikely to function as intended.

Option B, rpm -d, usually combined with other options like -q (query), is used for listing document files that are part of the package. It does not remove software. For example, rpm -qd mysoftware would list the documentation included within the "mysoftware" package.

Option C, rpm -q, stands for "query" and is used for querying information about installed packages. For instance, rpm -q mysoftware would check if the "mysoftware" package is installed. It does not remove software.

Therefore, only the -e option is designed explicitly for removing software packages. Package management is a cornerstone of system administration, especially crucial in cloud environments where consistent and automated deployments and removals are paramount. Incorrect package management can lead to instability, security vulnerabilities, and application malfunctions. Proper usage of rpm -e ensures controlled and clean software removal, minimizing the potential for negative impacts on the system. RPM helps maintain the desired state of virtual machines and container images in a cloud environment.

For further information, refer to the official RPM documentation and guides on Red Hat's website or similar resources. Also, research Linux package management best practices for server administration and cloud deployments.

RPM documentation
Red Hat Package Manager (RPM)

Question: 59

A Linux system fails to start and delivers the following error message:

```
Checking all file systems.
/dev/sda1 contains a file system with errors, check forced.
/dev/sda1: Inodes that were part of a corrupted orphan linked list found.
/dev/sda1: UNEXPECTED INCONSISTENCY;
```

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

A.fsck.ext4 /dev/sda1
B.partprobe /dev/sda1
C.fdisk /dev/sda1
D.mkfs.ext4 /dev/sda1

Answer: A

Explanation:

The fsck.ext4 command is used to check and repair file system errors on an ext4 file system. In this case, the

command is checking the file system located on /dev/sda1.The fsck.ext4 utility is used to detect and correct file system inconsistencies, such as corrupt or lost inodes, missing block groups, and other issues. It's typically run automatically by the operating system during system boot, or manually by the administrator when file system problems are suspected.By running the fsck.ext4 /dev/sda1 command, the administrator is checking the ext4 file system on the first partition of the first SATA disk (/dev/sda1) for any errors. If any problems are found, fsck.ext4 will attempt to repair them. It's important to note that running fsck.ext4 on a mounted file system can cause data loss, so it's recommended to run it on an unmounted file system or in a maintenance mode.

Question: 60

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

Explanation:

The /etc/login.defs file is used to set default system-wide settings for new users and groups. This file contains information about user and group creation, such as the default home directory, the default shell, the minimum and maximum UIDs and GIDs, and so on.If the administrator needs to set default values for new users and groups that are created, they would need to make the required changes in the /etc/login.defs file. For example, to set the default home directory for new users to /home/users/, the administrator would add the following line to the file:

Question: 61

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

Given the following analysis:

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

```
/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

Explanation:

File attributes are preventing file modification.

Question: 62

A Linux administrator needs to create a new cloud.cpio archive containing all the files from the current directory. Which of the following commands can help to accomplish this task?

A.ls | cpio -iv > cloud.epio
B.ls | cpio -iv < cloud.epio
C.ls | cpio -ov > cloud.cpio
D.ls cpio -ov < cloud.cpio

Answer: C

Explanation:

The correct command to create a new cloud.cpio archive containing all files from the current directory is ls | cpio -ov > cloud.cpio.

Here's why:

ls: This command lists the files and directories in the current directory, sending the output to standard output (stdout).

| (pipe): The pipe operator takes the stdout from the ls command and redirects it as standard input (stdin) to the cpio command. This allows cpio to receive the list of files.

cpio -ov: This invokes the cpio command with the following options:

-o (output): This tells cpio that we are creating an archive.
-v (verbose): This option displays a list of files as they are added to the archive (optional, but useful for monitoring progress).

> cloud.cpio: This redirects the standard output (stdout) from the cpio command to a file named cloud.cpio. This effectively creates the cpio archive.

Why other options are incorrect:

A. ls | cpio -iv > cloud.epio: The -i option tells cpio to extract files from an archive, not create one. Using > cloud.epio with extraction is incorrect; the redirection would overwrite a file named "cloud.epio" with possibly corrupted output from a failed extraction. Also, cloud.epio is a typo.

B. ls | cpio -iv < cloud.epio: Again, -i is for extraction. The < cloud.epio tries to provide the archive file as input, which is the opposite of what we want when creating an archive.

D. ls cpio -ov < cloud.cpio: This syntax is incorrect. It attempts to run ls and cpio as separate commands, and then redirect the content of cloud.cpio as stdin for the second command.

In summary, the ls | cpio -ov > cloud.cpio command chain correctly lists the files in the current directory, pipes that list to cpio to create an archive, and redirects the resulting archive data to the cloud.cpio file.

Further Reading:

cpio Manual:man cpio in your Linux terminal.
GNU cpio Documentation:https://www.gnu.org/software/cpio/manual/cpio.html

Question: 63

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

Answer: A

Explanation:

The correct answer is A, source ~/.bashrc. The .bashrc file is a shell script that Bash executes whenever a new interactive, non-login shell is started. When a user modifies .bashrc, those changes, including newly defined aliases, are not immediately available in the current shell. The source command, also represented by ., tells the current shell to read and execute the commands from the specified file in the current environment. This effectively reloads the .bashrc file, incorporating the new alias definition into the running shell session.

Without sourcing the file, the current shell session is unaware of the changes made to .bashrc. Typing the alias command will result in an "command not found" error, or unexpected behavior. read ~/.bashrc (option B) would simply read the content of the file but not execute it, rendering the changes ineffective. touch ~/.bashrc (option C) would only update the timestamp of the file, making it seem like it was modified, without implementing the changes to your current shell. echo ~/.bashrc (option D) would simply print the file path of the .bashrc file, this wouldn't solve the problem in any capacity. Using source ensures that the alias is loaded into the current shell's environment, enabling the administrator to use it immediately.

For further research, consider these resources:

GNU Bash Manual:https://www.gnu.org/software/bash/manual/bash.html (Specifically, search for ".bashrc" and "source".)
Linux Documentation Project:https://tldp.org/ (Search for "Bash scripting tutorial" or "introduction to the shell")

Question: 64

A junior systems administrator has just generated public and private authentication keys for passwordless login. Which of the following files will be moved to the remote servers?

A.id_dsa.pem
B.id_rsa
C.id_ecdsa
D.id_rsa.pub

Answer: D

Explanation:

The correct answer is id_rsa.pub because it represents the public key. Passwordless SSH authentication relies on a cryptographic key pair: a private key kept securely on the client machine and a corresponding public key deployed to the remote server. The client uses its private key to encrypt a challenge from the server, and the server uses the deployed public key to decrypt it. If the decryption is successful, authentication is granted without requiring a password.

id_rsa (option B) is the private key. Moving the private key to a remote server is a severe security vulnerability, as anyone with access to the server could then impersonate the user. Private keys must never be shared. The .pub extension clearly indicates the public key counterpart to the private key id_rsa. While id_dsa.pem (option A) and id_ecdsa (option C) also represent key files, the .pem extension with id_dsa signifies a Privacy Enhanced Mail format, though DSA (Digital Signature Algorithm) is less commonly used now. id_ecdsa indicates an Elliptic Curve Digital Signature Algorithm key, which is a valid type of key for SSH authentication, but what is transferred to the remote server in any of these situations is always the public part of the key. The public key (.pub) is safe to distribute because it cannot be used to decrypt data encrypted with the private key.
The public key allows the server to verify the identity of the client holding the corresponding private key. Transferring the public key to the ~/.ssh/authorized_keys file on the remote server enables passwordless login for the associated user.

Further resources:

SSH Key Management:https://www.ssh.com/academy/ssh/keygen
Understanding SSH Key Types:https://www.digitalocean.com/community/tutorials/understanding-ssh-key-pairs-openssh-keys
OpenSSH authorized_keys file format:https://man.openbsd.org/ssh/authorized_keys.5

Question: 65

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:

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: C

Explanation:

mount disk by-label.

https://unix.stackexchange.com/questions/644708/how-can-i-prevent-disk-uuid-mismatch-when-cloning-a-machine

Question: 66

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.

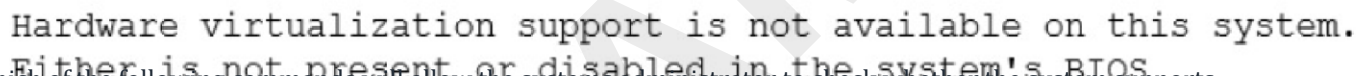

Which of the following commands will allow the systems administrator to check whether the system supports virtualization?

A.dmidecode -s system-version
B.lscpu
C.sysctl -a
D.cat /sys/device/system/cpu/possible

Answer: B

Explanation:

B. lscpuThe lscpu command is used to display information about the CPU architecture and the system topology on Linux systems. It provides information about the number of CPUs, cores, and threads, the

architecture, clock speed, and cache size. If the system supports virtualization, it should be reflected in the output of the lscpu command.The other options are not relevant for checking virtualization
support:dmidecode -s system-version is used to retrieve the version number of the system's BIOS or UEFI firmware.sysctl -a is used to display all system control parameters, including both kernel parameters and sysctl settings.cat /sys/device/system/cpu/possible does not exist, it should be /sys/devices/system/cpu/possible. This file provides information about the maximum number of CPUs that can be configured for the system, but does not indicate whether the system supports virtualization.

Question: 67

A Linux administrator created the directory /project/access2all. By creating this directory, the administrator is trying to avoid the deletion or modification of files from non-owners. Which of the following will accomplish this goal?

A.chmod +t /project/access2all
B.chmod +rws /project/access2all
C.chmod 2770 /project/access2all
D.chmod ugo+rwx /project/access2all

Answer: C

Explanation:

The correct answer is C, chmod 2770 /project/access2all. This command achieves the desired goal of preventing non-owners from deleting or modifying files created by others within the /project/access2all directory. Let's break down why.

The octal mode 2770 represents specific permissions. The '2' in the thousands place sets the SetGID (SGID) bit. The '7' in the hundreds place gives the owner full permissions (read, write, execute). The '7' in the tens place grants the group associated with the directory full permissions. Finally, the '0' in the ones place denies all permissions to others.

The crucial part is the SetGID bit. When SGID is set on a directory, any files created within that directory inherit the group ownership of the directory, rather than the group of the user who created the file. This is key because even if different users create files, they all belong to the same group.

Furthermore, because the 'others' permission is set to 0, users who are neither the owner nor members of the directory's group cannot modify or delete files within the directory. However, members of the group specified for the directory can create, read, write, and delete files.

Option A, chmod +t /project/access2all, sets the sticky bit. While the sticky bit prevents users from deleting or renaming files they don't own, even if they have write permissions on the directory, it doesn't affect modification. Moreover, everyone with write permissions can still create new files.

Option B, chmod +rws /project/access2all, is problematic. The 's' here attempts to set both SetUID and SetGID, but applying SetUID to a directory usually doesn't behave as expected. Besides, it still doesn't by itself guarantee that non-owners cannot modify existing files.

Option D, chmod ugo+rwx /project/access2all, grants everyone (user, group, and others) full permissions. This completely contradicts the requirement of preventing non-owners from deleting or modifying files.

In summary, chmod 2770 effectively restricts deletion/modification to the owner and members of the specified group, fulfilling the described security need.

For more detailed information, consider these resources:

chmod manual page:man chmod in a Linux terminal, or online (e.g., https://man7.org/linux/man-pages/man1/chmod.1.html)
Understanding Linux Permissions: Many online tutorials exist, such as those on Red Hat's developer site or through a simple search for "Linux file permissions." These tutorials often explain the octal notation used for chmod.

Question: 68

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.sysctl net.ipv4.ip_forward
B.sysctl -w net.ipv4.ip_forward=1
C.echo "net.ipv4.ip_forward=1" >> /etc/sysctl.conf
D.echo 1 > /proc/sys/net/ipv4/ip_forward
E.sysctl –p
F.echo "net.ipv6.conf.all.forwarding=l" >> /etc/sysctl.conf

Answer: BC

Explanation:

The correct answer is B and C. The goal is to enable IPv4 forwarding persistently.

Option B, sysctl -w net.ipv4.ip_forward=1, immediately enables IPv4 forwarding. The sysctl command allows modifying kernel parameters at runtime. The -w option makes the change. However, this change is not persistent across reboots.

Option C, echo "net.ipv4.ip_forward=1" >> /etc/sysctl.conf, ensures that the IPv4 forwarding setting is applied automatically at each boot. The /etc/sysctl.conf file is used to configure kernel parameters permanently. By appending net.ipv4.ip_forward=1 to this file, the system will set the net.ipv4.ip_forward parameter to 1 every time it starts.

Combining B and C guarantees that IPv4 forwarding is enabled immediately and persists across reboots. B applies the change right away and C ensures that the setting remains after rebooting the system.

Option A, sysctl net.ipv4.ip_forward, only displays the current value of the net.ipv4.ip_forward parameter, and doesn't change it. Option D, echo 1 > /proc/sys/net/ipv4/ip_forward, does enable forwarding, but this change is not persistent across reboots. Contents of /proc/sys are usually recreated at boot. Option E, sysctl -p, applies changes from /etc/sysctl.conf or a specified configuration file, but it's typically used after modifying /etc/sysctl.conf to activate the changes without rebooting and it doesn't by itself add the forwarding rule to the config file. Option F configures IPv6 forwarding, not IPv4.

Relevant Links:

sysctl: https://man7.org/linux/man-pages/man8/sysctl.8.html
/etc/sysctl.conf: https://www.man7.org/linux/man-pages/man5/sysctl.conf.5.html

Question: 69

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: C

Explanation:

The correct command to find and remove log files modified more than 180 days ago is find /var/log -type f -mtime +180 -exec rm \;. Let's break down why.

find /var/log: This initiates the find command and specifies the /var/log directory as the starting point for the search. Log files are typically stored in this directory on Linux systems.

-type f: This option filters the search results to include only files (regular files), ensuring that directories and other file types are excluded. This is important because we only want to remove log files, not directories.

-mtime +180: This is the critical part for the time-based filtering. -mtime refers to the file's modification time. +180 means that the command will only select files that were last modified more than 180 days ago.

-exec rm \;: This option executes the rm command (remove) on each file that matches the previous criteria. The is a placeholder that gets replaced by the name of each found file. The \; is required to terminate the -exec option.

Option A is incorrect because it uses -type d, which specifies directories, not files. This would attempt to remove directories in the /var/log directory which can lead to system instability.

Option B is incorrect due to using the invalid -modified flag, and the -rm command is not a standard option. The-exec rm option is necessary.

Option D is incorrect. -type c signifies character special file, not file. The -atime specifies the time of last access, and -remove is not a standard. The correct command for deleting files is rm.

Therefore, only option C correctly identifies the file type as a regular file, specifies the modification time criteria, and uses the correct rm command wrapped in -exec to remove the files.

Supporting links:

find command: https://man7.org/linux/man-pages/man1/find.1.html
/var/log directory: https://www.redhat.com/sysadmin/linux-log-files
rm command: https://man7.org/linux/man-pages/man1/rm.1.html

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 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
```

D.

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

Answer: C

Explanation:

C. echo 1 > /proc/sys/net/ipv4/ip_forwardiptables -t nat -A PREROUTING -o eth0 -j MASQUERADEExplanation:The first command (echo 1 > /proc/sys/net/ipv4/ip_forward) enables IP forwarding on the Linux server, allowing it to forward packets from one interface to another.The second command (iptables -t nat -A PREROUTING -o eth0 -j MASQUERADE) uses iptables to set up Network Address Translation (NAT), which will allow the Linux server to act as a router and forward traffic between the external network and the internal network. The "-o eth0" option specifies the outgoing interface, and the "-j MASQUERADE" option sets up MASQUERADE NAT, which dynamically assigns IP addresses to internal network clients as they make outbound connections.

Question: 71

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

```
Oct 20 03:45:50 hostnane kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx SRC=x.x.x.x DST=x.x.x.x LEN=1059 TOS=0x00
PREC=0x00 TTL=115 ID=31368 DF PROTO=TCP
SPT=17992 DPT=80 WINDOW=16477 RES=0x00 ACK PSH URGP=0
Oct 20 03:46:02 hostname kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx SRC=x.x.x.x DST=x.x.x.x LEN=52 TOS=0x00
PREC=0x00 TTL=52 ID=763 DF PROTO=TCP SPT=20229 DPT=22 WINDOW=15598 RES=0x00 ACK URGP=0
Oct 20 03:46:14 hostname kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx SRC=x.x.x.x DST=x.x.x.x LEN=324 TOS=0x00
PREC=0x00 TTL=49 ID=64245 PROTO=TCP SPT=47237 DPT=80 WINDOW=470 RES=0x00 ACK PSH URGP=0
Oct 20 03:46:26 hostname kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx SRC=x.x.x.x DST=x.x.x.x LEN=52 TOS=0x00
PREC=0x00 TTL=45 ID=2010 PROTO=TCP SPT=48322 DPT=80 WINDOW=380 RES=0x00 ACK URGP=0
```

Which of the following commands will remediate and help resolve the issue?

A.

```
IPtables -A FORWARD -i eth0 -p tcp --dport 80 -j ACCEPT
IPtables -A FORWARD -i eth0 -p tcp --dport 22 -j ACCEPT
```

B.

```
IPtables -A INPUT -i eth0 -p tcp --dport 80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --dport 22 -j ACCEPT
```

C.

```
IPtables -A INPUT -i eth0 -p tcp --sport 80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --sport 22 -j ACCEPT
```

D.

```
IPtables -A INPUT -i eth0 -p tcp --dport :80 -j ACCEPT
IPtables -A INPUT -i eth0 -p tcp --dport :22 -j ACCEPT
```

Answer: B

Explanation:

This iptables command is adding a new rule to the INPUT chain in the iptables firewall. The rule allows incoming traffic on interface "eth0" using the TCP protocol, destined for port 80, to be accepted. The "-A" option is used to append the rule to the end of the chain, and the "-j" option specifies the target action for the rule, in this case ACCEPT. This command is typically used to configure a firewall to allow specific types of traffic to enter the system.

Question: 72

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.systemctl restart sshd.service

D.chmod 600 mv .ssh/authorized_key

Answer: B

Explanation:

B. mv .ssh/authorized_key .ssh/authorized_keysThis changes the name of the file from authorized_key to authorized_keys, which is the conventional name for the file that contains authorized public keys for passwordless SSH connections. By convention, the ssh-keygen tool generates the authorized_keys file, and some SSH servers expect this exact name. Changing the name of the file to authorized_keys may resolve the issue and allow the passwordless connection to be established.

Question: 73

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: C

Explanation:

The file is owned by the "root" user (the owner) and belongs to the "root" group (the group). In many cases, system files should not be owned by the "root" user, especially if they are related to a specific service or application. It's generally recommended to have service-specific users and groups to manage permissions and access to files.

Question: 74

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

Explanation:

Here's a detailed justification for why option D is the correct answer and why the others are incorrect:

The task is to set the sticky bit on a directory named devops that already has file permissions of 755. The sticky bit, when applied to a directory, prevents users from deleting or renaming files within that directory unless they own the file, own the directory, or are the root user. This is particularly useful in shared directories like /tmp.

Why option D chmod 1755 devops is correct: The chmod command is used to change file permissions. To set the sticky bit numerically, you prefix the existing permission mode with a 1. The existing permission 755 already provides read, write, and execute permissions to the owner, read and execute permissions to the group, and read and execute permissions to others. Adding the 1 (representing the sticky bit) results in 1755.

This command correctly adds the sticky bit while preserving the existing permissions. The command chmod 1755 devops effectively sets the sticky bit on the devops directory, restricting file deletion or renaming within the directory to authorized users.

Why option A chown -s 755 devops is incorrect: The chown command is used to change the ownership of a file or directory. The -s option isn't a standard or widely recognized option for chown. Even if it existed, chown does not control permissions; chmod does. Attempting to use chown to modify permissions is fundamentally incorrect.

Why option B chown 1755 devops is incorrect: Again, chown is for changing ownership, not permissions. Even if chown did accept numerical permissions, you wouldn't set a file's owner to a numerical value like 1755; you would specify a user or user:group combination.

Why option C chmod -s 755 devops is incorrect: While chmod is the correct tool to modify permissions, the -s option doesn't directly set the sticky bit in the intended way using numerical mode. The -s option is associated with setting the setuid or setgid bits, depending on the context, and not for setting the sticky bit. Furthermore, using -s in this context would likely result in an error or unexpected behavior. The correct numerical prefix to set the sticky bit is 1.

In summary, chmod 1755 devops is the only command that correctly uses chmod with the proper numerical representation (1 as a prefix) to set the sticky bit while preserving the directory's existing 755 permissions.

Here are some authoritative links for further research:

chmod:https://man7.org/linux/man-pages/man1/chmod.1.html
Sticky Bit:https://en.wikipedia.org/wiki/Sticky_bit
File Permissions:https://www.redhat.com/sysadmin/linux-file-permissions

Question: 75

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: C

Explanation:

The correct answer is C because the administrator fixed the problem temporarily but didn't make the change permanent. The systemctl isolate graphical.target command only changes the current target. It switches the running system to the graphical target for that specific boot session. Upon reboot, the system reverts to its default target, which, in this scenario, is likely a non-graphical target like multi-user.target. To permanently set the graphical interface as the default, the administrator needs to use the systemctl set-default graphical.target command. This creates a symbolic link from /etc/systemd/system/default.target to /lib/systemd/system/graphical.target, instructing systemd to boot into the graphical target by default on every subsequent boot. Option A is incorrect because systemctl reboot is the proper way to reboot a systemd-based Linux system. Option B is incorrect because basic.target is a minimal target with very few services running, certainly not a desired state for general server operation; it would result in even fewer services running than

the non-GUI terminal. Option D is incorrect because the shutdown method (or lack thereof) doesn't persist the startup target configuration. The target is specified by the systemd configuration, not influenced by whether the server was shut down cleanly. The fundamental problem is the administrator's failure to persistently configure the system to boot into the graphical target.

For further information, consult the following resources:

systemd Documentation:https://www.freedesktop.org/wiki/Software/systemd/
Understanding Systemd Units and Targets:
https://www.digitalocean.com/community/tutorials/understanding-systemd-units-and-unit-files **Arch Linux Wiki on Systemd:**https://wiki.archlinux.org/title/systemd - A comprehensive and detailed resource.

Question: 76

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: C

Explanation:

C. Something is starving the server resources.The message "out of Memory: Killed process 1234 (mysqld)" indicates that the process mysqld was terminated by the Linux kernel due to a lack of available memory. This suggests that something is consuming all of the available memory on the server and preventing mysqld from functioning properly. This is an indication of resource starvation, which can cause various issues, including unexpected connection closures.Option A "The process mysqld is using too many semaphores" is not relevant to the issue described in the log message. Semaphores are used to synchronize processes, but they do not have an impact on memory usage.

Question: 77

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

A.$RHOST

B.SETENV

C.$SHELL

D.$DISPLAY

Answer: D

Explanation:

The correct answer is **D. $DISPLAY**.

Here's a detailed justification:

X11 forwarding allows you to run graphical applications on a remote Linux server and display them on your local machine. This is crucial when an application requires a graphical user interface (GUI) for installation or operation, and you're accessing the server via SSH from a machine that has a graphical environment.

The $DISPLAY environment variable is fundamental to this process. It tells the application which X server (the software responsible for drawing the GUI) to connect to for displaying its windows. When X11 forwarding is enabled in SSH, the SSH client sets up a tunnel between your local machine and the remote server. The server then creates a virtual X server (usually on a display number like :10.0 or similar) that is linked to this tunnel.

The $DISPLAY variable on the remote server must be set to point to this virtual X server. The SSH server typically handles this automatically when X11 forwarding is enabled. However, in some cases or when troubleshooting, you might need to verify or manually set it. If the $DISPLAY variable is incorrect or missing, the application won't know where to display its GUI, resulting in errors or a non-functional application. The SSH client on your local machine listens for connections on a local X server. When the remote application connects to the X server specified by the $DISPLAY variable, the SSH tunnel securely forwards the X11 protocol data to your local machine, allowing the GUI to be displayed.

Options A, B, and C are incorrect because they serve different purposes. $RHOST typically represents the remote host's name or IP address but is not directly related to X11 forwarding. SETENV is not a standard environment variable in Linux or shell scripting related to display functionality. $SHELL indicates the user's default shell (e.g., /bin/bash) and has no bearing on X11 display routing.

Here are some authoritative links for further research:

SSH X11 Forwarding:https://www.ssh.com/academy/ssh/x11-forwarding
Understanding X11:https://en.wikipedia.org/wiki/X_Window_System
DISPLAY environment variable:https://unix.stackexchange.com/questions/9786/what-is-the-display-environment-variable-used-for

Question: 78

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: BC

Explanation:

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.

Question: 79 A Linux administrator needs to correct the permissions of a log file on the server. Which of the following commands should be used to set filename.log permissions to -rwxr—r--. ?

A.chmod 755 filename.log
B.chmod 640 filename.log
C.chmod 740 filename.log
D.chmod 744 filename.log

Answer: D

Explanation:

The correct command is chmod 744 filename.log because it accurately translates the desired permissions (-rwxr--r--) into its numerical representation. Let's break down why.

The chmod command is used to change file permissions in Linux. The permission string -rwxr--r-- is symbolic notation. In the numerical (octal) representation, each digit represents permissions for a specific user category: owner, group, and others.

The first digit represents the owner's permissions. rwx (read, write, execute) translates to 4+2+1 = 7. The second digit represents the group's permissions. r-- (read only) translates to 4+0+0 = 4.
The third digit represents the permissions for others. r-- (read only) translates to 4+0+0 = 4.

Therefore, -rwxr--r-- is numerically represented as 744.

Option A (chmod 755 filename.log) would set the permissions to -rwxr-xr-x, granting execute permissions to the group and others, which is not the required permission.

Option B (chmod 640 filename.log) would set the permissions to -rw-r----- granting write access to the user only, which is incorrect.

Option C (chmod 740 filename.log) would set the permissions to -rwxr----- not granting read access for others, which is also incorrect.

The command chmod 744 filename.log correctly sets the file permissions for the owner (user) to read, write, and execute, and read-only access for the group and others. This aligns precisely with the desired permission string of -rwxr--r--. This is a fundamental concept of Linux system administration, crucial for managing file access and security.

Further research can be conducted at:

GNU chmod: https://www.gnu.org/software/coreutils/manual/html_node/chmod-invocation.html Understanding Linux File Permissions: https://www.linux.com/training-tutorials/understanding-linux-file-permissions/

Question: 80

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: D

Explanation:

The correct answer is **D. chage -E -1 accountname**. Here's why:

chage (Change User Password Expiration Information): The chage command is specifically designed for modifying user password expiration information. This includes setting expiration dates, last password change dates, and inactivity periods.

-E (Expire Date): The -E option in the chage command is used to set the expiration date for a user account.

-1 (Never): When used with the -E option, the value -1 signifies that the account should never expire. This effectively removes any expiration date that was previously set.

accountname: This refers to the username of the account whose expiration date you wish to remove.

Therefore, chage -E -1 accountname instructs the system to change the expiration date (-E) of the specified accountname to never expire (-1), effectively removing the existing expiration date.

Why the other options are incorrect:

A. chgrp system accountname: The chgrp command is used to change the group ownership of a file or directory. It has nothing to do with user account expiration dates.

B. passwd –s accountname: The passwd command with the -s option is used to display password status information for an account. It doesn't modify the expiration date.

C. chmod -G system account name: The chmod command is used to change file or directory permissions. The -G option specifies the group ID of the file. It's not related to user account expiration.

Authoritative Links:

chage command:man chage in a Linux terminal or https://man7.org/linux/man-pages/man1/chage.1.html **passwd command:**man passwd in a Linux terminal or https://man7.org/linux/man-pages/man1/passwd.1.html **chmod command:**man chmod in a Linux terminal or https://man7.org/linux/man-pages/man1/chmod.1.html **chgrp command:**man chgrp in a Linux terminal or https://man7.org/linux/man-pages/man1/chgrp.1.html