



- Expert Verified, Online, **Free**.



CERTIFICATION TEST

- CertificationTest.net - Cheap & Quality Resources With Best Support

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`

Suggested Answer: A

🗲️ 👤 **BENTECHY** 4 months, 1 week ago

Selected Answer: A

`rpm` and `grep` commands allow the Admin to establish the installed kernel packages while the `uname` with the option `(-a)` displays information about the current version of kernel on `/boot/kernel` file.

upvoted 3 times

🗲️ 👤 **ericsrz** 6 months, 4 weeks ago

A. query and grep kernel; information

upvoted 2 times

🗲️ 👤 **linux_admin** 1 year, 4 months ago

A. `rpm -qa | grep kernel; uname -a`

The administrator can use the command "`rpm -qa | grep kernel`" to identify the installed kernel packages on the system, and "`uname -a`" to display information about the current running kernel. The output of these commands can be used to determine the correct version of the `/boot/vmlinuz` file that needs to be restored.

upvoted 4 times

🗲️ 👤 **bjornborg** 1 year, 8 months ago

A correct....see what's installed, and see what's running

upvoted 2 times

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

Suggested Answer: D

🗲️ 👤 **linux_admin** Highly Voted 10 months, 3 weeks ago
D. /etc/ssh/sshd_config

The file the engineer needs to modify is "/etc/ssh/sshd_config". This is the configuration file for the SSH server and it controls various parameters for the server, including the port number used for secure remote login. To change the port number from 22 to 49000, the engineer should locate the line that starts with "Port 22" and change it to "Port 49000". After making the change, the engineer should save the file and restart the SSH service for the changes to take effect.

upvoted 7 times

🗲️ 👤 **noxkrugger** Most Recent 5 days, 7 hours ago
Selected Answer: D
Only modify in the /etc/ssh
upvoted 1 times

🗲️ 👤 **Nvoid** 1 year, 1 month ago
was on the test.
upvoted 3 times

🗲️ 👤 **TheRealManish** 1 year, 2 months ago
D is correct
upvoted 1 times

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

Suggested Answer: B

🗨️ 👤 **linux_admin** 10 months, 3 weeks ago

B. git pull

The administrator should use the "git pull" command to synchronize the local repository with the main repository. The "git pull" command fetches the latest changes from the remote repository and merges them with the local repository. This will ensure that the local repository has the latest version of the code, including the new file that was added to the main repository. Before running the "git pull" command, the administrator should make sure that they are in the correct branch and that their local repository is in a clean state (i.e., no changes have been made to the local repository since the last time it was synchronized with the main repository).

upvoted 2 times

🗨️ 👤 **Huckleberry** 11 months, 2 weeks ago

Pulling is the automated version of git fetch, it downloads a branch from a remote repository, then immediately merges it into the current branch.

upvoted 2 times

🗨️ 👤 **TheRealManish** 1 year, 2 months ago

GIT PULL is a mixture of fetch and merge.. It is the only answer that is relevant here

upvoted 1 times

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 tcp --dport 81 -j DNAT --to-destination 192.0.2.25:3129`
- C. `iptables -t nat -I PREROUTING -p tcp --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`

Suggested Answer: D

Community vote distribution

D (100%)

linux_admin **Highly Voted** 1 year, 4 months ago

The command "`iptables -t nat -A PREROUTING -p tcp --dport 80 -j DNAT --to-destination 192.0.2.25:3128`" will temporarily redirect all HTTP traffic to the new proxy server at IP address 192.0.2.25 and port 3128. The options used in this command are:

- t nat: Specifies the table (nat) to be used for this operation
- A PREROUTING: Specifies that the rule should be added to the PREROUTING chain, which is used to manipulate incoming packets before they are processed by the routing decision
- p tcp: Specifies that the rule should apply to TCP packets
- dport 80: Specifies that the rule should apply to incoming packets destined for port 80 (the default HTTP port)
- j DNAT: Specifies that the target of the rule should be DNAT (Destination NAT), which rewrites the destination address of a packet
- to-destination 192.0.2.25:3128: Specifies the destination address and port to which incoming packets should be redirected

Note: After executing this command, the administrator should make sure to save the iptables configuration, so that it will persist after a reboot.

upvoted 5 times

Alizadeh **Most Recent** 10 months, 2 weeks ago

Selected Answer: D

The correct answer is D. `iptables -t nat -A PREROUTING -p tcp --dport 80 -j DNAT --to-destination 192.0.2.25:3128`.

upvoted 1 times

Huckleberry 1 year, 5 months ago

- A --append
- D --delete
- t --tables
- nat PREROUTING(for altering packets as soon as they come in)

upvoted 2 times

bjornborg 1 year, 8 months ago

D is correct. Learn more here

<https://sourcedaddy.com/networking/running-public-services-private-ip-addresses.html>

upvoted 1 times

Nvoid 1 year, 7 months ago

was on the test and i choose D.

upvoted 2 times

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 eth0 -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`

Suggested Answer: D

Community vote distribution

D (100%)

linux_admin **Highly Voted** 1 year, 4 months ago

D. `ip route add 10.0.213.5/32 via 10.0.5.1 dev eth0`

The administrator should run the following command to implement the persistent, static route:

`"ip route add 10.0.213.5/32 via 10.0.5.1 dev eth0"`

This command adds a static route to the routing table, which specifies that packets sent over the interface "eth0" to the destination address "10.0.213.5/32" should be routed via the gateway "10.0.5.1". The "ip" command is used to manage the Linux IP routing table, and the "route add" option is used to add a new route. The "dev" option specifies the interface over which the packets should be sent. To make the route persistent across reboots, the administrator can add the same command to the appropriate configuration file (e.g., /etc/rc.local).

upvoted 7 times

Alizadeh **Most Recent** 10 months, 2 weeks ago

Selected Answer: D

The correct answer is D. `ip route add 10.0.213.5/32 via 10.0.5.1 dev eth0`.

upvoted 1 times

TheRealManish 1 year, 8 months ago

D seems to be correct. I tested it locally and it worked.

upvoted 1 times

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

Suggested Answer: A

🗨️ **Veteran903** Highly Voted 1 year, 1 month ago

please correct the mistake, right after the hashtag should be a bang(!) and not a number 1
upvoted 16 times

🗨️ **linux_admin** Highly Voted 10 months, 3 weeks ago

he command "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.

upvoted 9 times

🗨️ **MrJ_** Most Recent 10 months, 3 weeks ago

Answer must be A -f file True if file exists and is a regular file THEN

upvoted 3 times

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.

Snippets

done	while	xz
pgrep	locate	for
sed	filename	until
tar	then	\$log
egrep	"\$log.[1-6]\$"	"log.[1-6]@"
"\$1"	repeat	in
/tmp/tmpfile	/var/log	zip
gzip	rar	awk
"\$6"		

```
#!/bin/bash

#name: script.sh

find /var/log -type f -maxdepth 1 | grep [?] > /tmp/tmpfile

[?] filename [?] $(cat [?])

do

    [?] $filename

    [?]
```

Suggested Answer:

"log.[1-6]@"

for

in

/tmp/tmpfile

gzip

done

linux_admin Highly Voted 1 year, 10 months ago

The command "grep ".log[1-6]@" is a regular expression that is used to search for files with a ".log1" to ".log6" extension. The grep command is used to search for text patterns in files, and the regular expression provided as an argument is used to specify the pattern to search for.

The regular expression ".log[1-6]@" matches any string that ends with ".log1", ".log2", ".log3", ".log4", ".log5", or ".log6". The . character matches any single character, the log matches the characters "log", and the [1-6] matches any single digit in the range 1 to 6. The \$ symbol matches the end of the string, ensuring that the match only occurs if the ".log[1-6]@" pattern is at the end of the string.

upvoted 5 times

TheMichael Most Recent 4 months, 1 week ago

So the question is asking to compress individual rotating log files without recursion which sounds confusing, but it really just means it wants you to compress any file found in the log directory and no other directories within it.



upvoted 1 times

TheMichael 4 months, 1 week ago

They then write -maxdepth 1 which fulfills the "without recursion" portion of the question for you, so you can really just focus on just compressing the files in the /var/ log directory. So if you search (grep) for "\$log.[1-6]@" you are searching for a variable called log... which isn't a variable that exists in our code. If you search for "log.[1-6]@" you are only searching for the files that end in log.1 log.2 log.3... to log.6. Obviously you want to search for any file in the directory, so we use "\$1" to be able to search for the first thing we write after executing the code. "\$6" would be searching

for the 6th word we input after putting in the code, which doesn't really make sense to do when we can just search for the first term after typing the code in to execute it:

upvoted 1 times

  **TheMichael** 4 months, 1 week ago

ex:

script.sh yes no maybe so today tomorrow

would search for the first term out of the options yes no maybe so today tomorrow. Thus we would search for yes if this were to be entered exactly like this. "\$6" would search for the 6th term so that would be searching for tomorrow if entered this way.

the "for filename in" part is saying that we are creating a variable called filename, and FOR the variable we are putting IN this word:

we then write a cat command to make the filename variable be every word in the /tmp/ tempfile that we wrote the output of our grep search earlier to.

upvoted 1 times

  **TheMichael** 4 months, 1 week ago

The final do gzip command then just reads the variable we created called filename and compresses it for each term that it changes to from our for filename in command. This creates a single zipped file for each term found that matches our search within /var/ log, which is essentially known as compressing rotated log files without recursion aka we gzipped files we searched for only in the /var/ log directory and no additional directories within that.

since the wording could be interpreted as finding specifically .log files within the /var/ log directory, the "\$1" command allows us to specifically search for only those files if we so choose. The log.[1-6] option restricts us to only the files that end in log.[1-6] which wouldn't allow us to compress a file if it was log.11 or log.7, which doesn't make sense to do as the question doesn't ask us to specifically look for log files between 1 and 6, but rather log files in general.

upvoted 1 times

  **TheMichael** 4 months, 1 week ago

```
#!/bin/bash
```

```
#name: script.sh
```

```
find /var/log -type f -maxdepth 1 | grep "$1" > /tmp/tempfile
```

```
for filename in $(cat /tmp/tempfile)
```

```
do
```

```
gzip $filename
```

```
done
```

Is the final answer.

upvoted 2 times

  **DRVision** 1 year ago

```
#!/bin/bash
```

```
#name: script.sh
```

```
find /var/log -type f -maxdepth 1 | grep "$1" > /tmp/tempfile
```

```
for filename in $(cat /tmp/tempfile)
```

```
do
```

```
gzip $filename
```

```
done
```

The \$1 takes the first command line argument so it allows you to search what you're actually looking for through the logs, i.e. "related files".

Otherwise, even if the syntax was a typo for "log.[1-6]\$", because it should be ".log[1-6]\$", all you're doing is searching for logs, and it doesn't allow you to specify what you're looking

upvoted 2 times

  **DRVision** 1 year ago

thus allowing you to run the script and provide the search term into \$1 via command line

upvoted 2 times

🗨️ **Damon54** 1 year, 3 months ago

```
#!/bin/bash
```

```
#name: script.sh
```

```
find /var/log -type f -maxdepth 1 | grep "$1" > /tmp/tempfile
```

```
for filename in $(cat /tmp/tempfile)
```

```
do
```

```
gzip $filename
```

```
done
```

upvoted 2 times

🗨️ **Alizadeh** 1 year, 4 months ago

```
#!/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
```

upvoted 1 times

🗨️ **linux_admin** 1 year, 10 months ago

```
#!/bin/bash
```

```
# Find files in /var/log directory without recursion
```

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

```
# Loop through each file in /tmp/tempfile
```

```
for filename in $(cat /tmp/tempfile)
```

```
do
```

```
# Compress the file using gzip
```

```
gzip $filename
```

```
done
```

upvoted 2 times

🗨️ **linux_admin** 1 year, 10 months ago

1. Finds files in the /var/log directory without recursion: The find command is used to search for files in the /var/log directory. The -type f option limits the search to regular files, and the -maxdepth 1 option specifies that the search should not go deeper than the first level of subdirectories. The grep ".log.[1-6]\$" option filters the results to only include files with a ".log.1" to ".log.6" extension. The filtered results are redirected to /tmp/tempfile for processing.

2. Loops through each file in /tmp/tempfile: The for loop is used to iterate through each file listed in /tmp/tempfile. The \$filename variable holds the name of the current file in each iteration of the loop.

3. Compresses each file using gzip: The gzip command is used to compress each file. The file name is passed as an argument to the gzip command, and the compressed file will have a .gz extension added to its original name.

upvoted 3 times

🗨️ **linux_admin** 1 year, 10 months ago

The script will compress all files with the ".log.1" to ".log.6" extensions in the /var/log directory, and it will not go deeper than the first level of subdirectories. The compressed files will be stored in the same directory as the original files.

upvoted 1 times

🗨️ **Mr_Marcus** 1 year, 10 months ago

Which is all fabulous, except you missed several critical points in your analysis. First, the question does not ask for files with "log" in the name (it doesn't even ask for log files - it says "related files"). Second, some of the files in /var/log do not have log or .log anywhere in the filename (e.g. messages, secure, spooler, etc.) and would be missed using your syntax. Third, ".log[1-6]\$", with a leading dot (.), is not even an option to select. Your choices are "log.[1-6]\$" or "\$log.[1-6]\$", without a leading dot.

Your version will (likely) create an empty tempfile and compress nothing. And, yes, I've tested your solution (and mine).

The correct answer is:

```
#!/bin/bash
```

```
#name: script.sh
```

```
find /var/log -type f -maxdepth 1 | grep "$1" > /tmp/tempfile
```


```
for filename in $(cat /tmp/tempfile)
```

```
do
```

```
gzip $filename
```

```
done
```

upvoted 8 times

  **nabalauski** 1 year ago

after doing some research myself I can say with confidence that Mr_Marcus is correct here

upvoted 2 times

  **Lwarder1** 1 year, 11 months ago

Why is it not `grep "$1" > /tmp/tempfile`

upvoted 1 times

  **TheRealManish** 2 years ago

I can't figure out why we are grepping for `"log[1-6]$"` i dont see anything in the qusetion about 1-6.. but if we grep for `"$log"` we get all of the log files..

upvoted 1 times

  **CodeMaestro** 1 year, 8 months ago

For that we have a `$log` and not `log`

upvoted 1 times



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

```
resource "aws_instance" "ec2_instance" {  
  
    ami                        = data.aws_ami.vendor-Linux-2.id  
    associate_public_ip_address = true  
    count                     = 3  
    instance_type             = "instance_type"  
    vpc_security_group_ids    = [aws_security_group.allow_ssh.id]  
    key_name                   = aws_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

Suggested Answer: D

 **linux_admin**  1 year, 4 months ago

Here is an example code for deploying cloud-based servers using Terraform:

```
provider "aws" {  
    region = "us-west-2"  
}  
  
resource "aws_instance" "example" {  
    ami = "ami-0c55b159cbf0e1f0"  
    instance_type = "t2.micro"  
  
    tags = {  
        Name = "example-instance"  
    }  
}  
  
upvoted 6 times
```

 **linux_admin** 1 year, 4 months ago

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.

upvoted 8 times

  **JSHack** Most Recent 6 months, 3 weeks ago

Selected Answer: D

This is code used by Terraform.

upvoted 2 times

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

Suggested Answer: A

Community vote distribution

A (100%)

linux_admin **Highly Voted** 1 year, 4 months ago

A. LDAP (Lightweight Directory Access Protocol) can be used as a central repository of Linux users and groups. LDAP is a widely used protocol for accessing directory services and is commonly used to store user and group information in a central repository. This central repository can be used to manage user and group information for multiple systems, making it easier to manage and maintain user and group information across an organization.

upvoted 8 times

JSHack **Most Recent** 6 months, 3 weeks ago

Selected Answer: A

LDAP (Lightweight Directory Access Protocol): A protocol used to access and maintain distributed directory information services over a network. It is commonly used for authentication and managing user information in enterprise environments.

MFA (Multi-factor Authentication): A security process in which a user is required to provide two or more verification factors to gain access to a resource, such as a website or application. These factors typically include something the user knows (password), something the user has (a mobile phone or token), and something the user is (fingerprint or other biometrics).

SSO (Single Sign-On): A user authentication process that allows a user to access multiple applications with a single set of login credentials. This reduces the need to remember multiple passwords and improves user convenience.

PAM (Pluggable Authentication Module): A framework that allows system administrators to configure authentication methods for applications on Unix-like systems. It enables the integration of different authentication technologies (like password-based, biometrics, or MFA) without changing the applications themselves.

upvoted 2 times

Alizadeh 10 months, 2 weeks ago

Selected Answer: A

The answer is A. LDAP.

upvoted 2 times

lizb7223 1 year, 5 months ago

A

LDAP= Lightweight Director Access Protocol

MFA= Multi-factor Authentication

SSO= Single Sign On

PAM= Pluggable Authentication Module

upvoted 4 times

bjornborg 1 year, 8 months ago

Selected Answer: A

lightweight directory access protocol

others are multi-factor authentication, single sign on, and pluggable authentication modules...they can use ldap for authenticating

upvoted 2 times

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

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

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

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

Suggested Answer: D

linux_admin Highly Voted 1 year, 4 months ago

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.

upvoted 8 times

JSHack Most Recent 6 months, 3 weeks ago

Selected Answer: D

D. The network interface cable is not connected to a switch.

"state DOWN" indicates there is a problem with the network connection. Either the NAT is disconnected or disabled.

upvoted 1 times

Timebear 1 year, 6 months ago

The fact that it says 'no carrier' should be dead giveaway to there being no cable.


upvoted 4 times

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`

Suggested Answer: A

Community vote distribution



🗳️ 👤 **JRS99** 6 months, 2 weeks ago

Selected Answer: A

A. HOST:CONTAINER, 443 on the host and 8443 443:8443
upvoted 2 times

🗳️ 👤 **Mike313** 10 months, 1 week ago

A. -p HOST:CONTAINER. (Maps Port)
B Invalid as that is syntaxed as CONTAINER:HOST
C Invalid as it uses the -e option (ENV VARIABLE) which is not used
D Invalid for CONTAINER:HOST among other things.
upvoted 2 times

🗳️ 👤 **rfc_1918** 1 year, 5 months ago

A. If you're using docker compose the pattern is HOST:CONTAINER, i.e. 443 on the host and 8443 on the container would be 443:8443
upvoted 2 times

🗳️ 👤 **BryanSME** 1 year, 6 months ago

<https://www.howtoforge.com/getting-started-with-podman-manage-images-container-and-volumes/>
which is mentioned by Rob74613 also indicates host:container so A is correct
Manage Containers

At this stage, we learned about container image management. And the next step, we will learn how to create and manage containers with Podman.

To create a new container, we can use 'podman run' command as below.

```
podman run -d -p 8000:80 --name hakase-nginx docker.io/library/nginx
```

The command will create a new container named 'hakase-nginx' based on the nginx image and will expose the port 8000 on the host machine.

Details command options:

-d - keep the container running in the background and just print the container ID as a result.
-p 8000:80 - port mapping for container and the host system. Port 8000 on the host machine, and port 80 on the container.
--name hakase-nginx - specify the container name with 'hakase-nginx'.

Now display all running containers on the system.
upvoted 2 times

🗳️ 👤 **BryanSME** 1 year, 6 months ago

This is a useful article:
<https://docs.podman.io/en/latest/markdown/podman-run.1.html>
-p=[[ip:][hostPort:]containerPort[/protocol]]
this would make "A" the correct answer

upvoted 1 times

🗳️ 👤 **nabalauski** 1 year, 6 months ago

Selected Answer: B

answer is B.

The exposed port of an application can be mapped to a host port using the `-p` flag. For example, an `httpd` port 80 can be mapped to the host port 8080 using the following:

```
$ podman run -p 8080:80 -d -i -t fedora/httpd
```

upvoted 1 times

🗳️ 👤 **[Removed]** 1 year, 6 months ago

A should be correct

upvoted 1 times

🗳️ 👤 **DRVision** 1 year, 6 months ago

Selected Answer: B

To create a new container, we can use 'podman run' command as below.

```
podman run -d -p 8000:80 --name hakase-nginx docker.io/library/nginx
```

The command will create a new container named 'hakase-nginx' based on the `nginx` image and will expose the port 8000 on the host machine.

Details command options:

`-d` - keep the container running in the background and just print the container ID as a result.

`-p 8000:80` - port mapping for container and the host system. Port 8000 on the host machine, and port 80 on the container.

`--name hakase-nginx` - specify the container name with 'hakase-nginx'.

Now display all running containers on the system.

upvoted 1 times

🗳️ 👤 **funax** 1 year, 7 months ago

Selected Answer: A

`<host_port>:<container_port>`.

upvoted 2 times

🗳️ 👤 **clean_it_up_janny** 1 year, 11 months ago

Selected Answer: B

concur with rob

upvoted 1 times

🗳️ 👤 **Rob74613** 2 years ago

Answer is B

Source: <https://www.howtoforge.com/getting-started-with-podman-manage-images-container-and-volumes/>

look at section 3: Manage Containers

upvoted 3 times

🗳️ 👤 **mrtwister76** 2 years, 1 month ago

Definitely A

upvoted 1 times

🗳️ 👤 **tutita** 2 years, 1 month ago

Selected Answer: A

I checked the documentation and first is host port and then the container port, so option A should be the right one. since you are exposing the port 443 on the host to the port 8443 on the container

upvoted 2 times

🗳️ 👤 **Rob74613** 2 years, 1 month ago

Selected Answer: B

B is correct

`-p 8443:443`: This option specifies the port mapping. It maps port 8443 of the host machine to port 443 of the container.

upvoted 1 times

🗨️ 👤 **Aamm033** 2 years, 3 months ago

Selected Answer: A

Ans is A.

upvoted 1 times

🗨️ 👤 **kloug** 2 years, 3 months ago

bbbbbbbbb

upvoted 1 times

🗨️ 👤 **nixonbii** 2 years, 4 months ago

According to the Podman documentation, the host machine port is the first parameter in the command when used with the -p option:

`-p=[[ip:][hostPort]:]containerPort[/protocol]`

Under this construction, answer A would be correct. I know that Podman can be seamlessly aliased to Docker in Ubuntu but I am not sure if Docker uses a different construction. I also think that the fact that Podman is a daemonless application influences how it interfaces with the O/S ports and protocols.

upvoted 2 times


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`

Suggested Answer: B

Community vote distribution

B (100%)

 **Aamm033** Highly Voted 1 year, 2 months ago

Selected Answer: B

B is correct... exec option is for running containers -t for TTY and -i interactive
upvoted 5 times

 **linux_admin** Most Recent 10 months, 3 weeks ago

B. `podman exec -ti app /bin/sh`

The podman exec command allows a Linux administrator to enter a running container and perform operations inside the container. The -ti options specify that the administrator wants to run the command interactively with a TTY and with STDIN attached.

The app in the command specifies the name or ID of the running container that the administrator wants to enter. The /bin/sh at the end of the command specifies the shell to use inside the container.
upvoted 4 times

 **linux_admin** 10 months, 3 weeks ago

In this case, the administrator wants to enter the container to analyze logs that are stored inside. To accomplish this, the administrator should run the following command:

```
podman exec -ti app /bin/sh
```

This command will enter the running container named app and open a shell inside the container. The administrator can then use standard Linux commands to analyze the logs and diagnose the issue with the failing application
upvoted 1 times

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

Suggested Answer: C

Community vote distribution

C (100%)

linux_admin **Highly Voted** 1 year, 4 months ago

C. dd if=/dev/sdc1 of=/dev/sdd1

The dd command is a low-level utility that can be used to copy data from one location to another. In this case, the administrator wants to clone the partition /dev/sdc1 to /dev/sdd1. To accomplish this, the administrator should run the following command:

```
dd if=/dev/sdc1 of=/dev/sdd1
```

The if option specifies the input file, which is /dev/sdc1 in this case. The of option specifies the output file, which is /dev/sdd1. The dd command will read the data from /dev/sdc1 and write it to /dev/sdd1, effectively cloning the partition.

Note that the dd command is a very powerful tool and can cause data loss if used improperly. The administrator should make sure to backup important data before using the dd command, and be very careful when specifying the input and output files.

upvoted 5 times

Alizadeh **Most Recent** 10 months, 2 weeks ago

Selected Answer: C

The correct answer is C. dd if=/dev/sdc1 of=/dev/sdd1.

upvoted 2 times

Aamm033 1 year, 8 months ago

Selected Answer: C

C is correct

upvoted 3 times

bjornborg 1 year, 8 months ago

Selected Answer: C

dd is the only command that can do this from the above

if=input-file of=output file

upvoted 2 times

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.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **linux_admin** 10 months, 3 weeks ago

The /bin/false shell is a special shell that returns a non-zero exit status, effectively denying the user access to the system. This is often used for system accounts that do not need to interact with the system, or for temporarily disabling accounts.

upvoted 4 times

🗳️ 👤 **linux_admin** 10 months, 3 weeks ago

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

upvoted 4 times

🗳️ 👤 **Aamm033** 1 year, 2 months ago

Selected Answer: D

D correct

upvoted 1 times

🗳️ 👤 **bjornborg** 1 year, 2 months ago

option 1, last column after : says /bin/false ... that usually says /bin/bash for the user's shell ... /bin/false would deny user being able to login

upvoted 1 times

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

Suggested Answer: *DF*

🗲️ 👤 **linux_admin** Highly Voted 10 months, 3 weeks ago

- D. ssh-copy-id
- F. scp

The ssh-copy-id and scp commands can be used to copy an SSH key file to remote servers.

The ssh-copy-id command is used to install the public key of an SSH key pair to the authorized_keys file of a remote server. This allows the user to log into the remote server without a password, using the private key of the key pair.

upvoted 7 times

🗲️ 👤 **linux_admin** 10 months, 3 weeks ago

For example:

```
ssh-copy-id user@remote_server
```

The scp command is used to securely copy files between systems over an SSH connection. The command can be used to copy an SSH key file to a remote server.

For example:

```
scp ~/.ssh/id_rsa.pub user@remote_server:/tmp/id_rsa.pub
```

upvoted 7 times

🗲️ 👤 **linux_admin** 10 months, 3 weeks ago

In this example, the public key of the SSH key pair (~/.ssh/id_rsa.pub) is copied to the remote server (remote_server) in the /tmp directory with the name id_rsa.pub. The administrator can then use the ssh-copy-id command to install the public key on the remote server.

upvoted 3 times

🗲️ 👤 **Deuteronomy** Most Recent 10 months, 3 weeks ago

D and F are correct. ssh-keygen command is used to generate a pair of keys.

upvoted 1 times

🗲️ 👤 **lzb7223** 11 months ago

BD

https://linuxhint.com/copy_ssh_keys/

upvoted 2 times

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`

Suggested Answer: B

🗲️ 👤 **linux_admin** Highly Voted 10 months, 3 weeks ago

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

The `sysctl` command is used to configure the Linux kernel parameters at runtime. To enable IPv4 packet forwarding, the administrator needs to set the value of the `net.ipv4.ip_forward` parameter to 1. This can be done with the following command:

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

This command sets the value of `net.ipv4.ip_forward` to 1, enabling IPv4 packet forwarding. The `-w` option specifies that the change should be written to the running system and not just displayed.

upvoted 8 times

🗲️ 👤 **linux_admin** 10 months, 3 weeks ago

Note that this change will only persist until the next reboot. To make the change persistent across reboots, the administrator needs to add the following line to the `/etc/sysctl.conf` file:

```
net.ipv4.ip_forward=1
```

After adding this line, the administrator can run the `sysctl -p` command to reload the `/etc/sysctl.conf` file and apply the changes.

upvoted 5 times

🗲️ 👤 **lizb7223** Most Recent 11 months ago

B

<https://linuxconfig.org/how-to-turn-on-off-ip-forwarding-in-linux>

upvoted 1 times

🗲️ 👤 **Hava_2013** 1 year, 1 month ago

<https://linuxconfig.org/how-to-turn-on-off-ip-forwarding-in-linux>

upvoted 3 times

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.

Suggested Answer: D

Community vote distribution

C (100%)

🗨️ 👤 **Veteran903** Highly Voted 2 years, 7 months ago

Hello all,

I respectfully disagree, answer is correct, its D, there is nothing wrong with the OnCalendar, you ALWAYS need to reload system-daemon after you change or update configurations, this is how you make the system aware of the changes, This will rerun all generators, reload all unit files, and recreate the entire dependency tree.

upvoted 5 times

🗨️ 👤 **Nvoid** 2 years, 7 months ago

They did reload system-daemon, by rebooting the server after the change. Read the question again.

upvoted 2 times

🗨️ 👤 **Veteran903** 2 years, 7 months ago

restarting the server DO NOT reload de demon sir, you must do it manually or using an script but NEVER by restarting the server....the answer is D, you all are not seeing through this question...if you make changes you MUST MANUALLY reload the demon to make the system aware of the change.

upvoted 3 times

🗨️ 👤 **mrtwister76** 2 years, 1 month ago

When you restart a Linux server, the init system (such as systemd) automatically reloads the unit files, so you don't need to explicitly run the command `systemctl daemon-reload` to reread the unit files.

The `systemctl daemon-reload` command is typically used when you make changes to the unit files manually and want to inform the init system to reload the updated configuration without restarting the entire server. This command tells systemd to reparse the unit files and update its internal configuration.

However, during a server restart, systemd performs a complete reload of the unit files as part of the initialization process. It reads the unit files from their defined locations and applies the configuration without the need for an explicit `daemon-reload` command.

So, when you restart the server, systemd will automatically reload the unit files, ensuring that the updated configuration is applied, and the services are started based on the latest settings.

The correct answer is C. :)

upvoted 4 times

🗨️ 👤 **Mike313** Most Recent 10 months, 1 week ago

Selected Answer: C

The answer is C

It CANNOT be A, B, simply because the job would not execute AT ALL if there was an issue with the service.

It also CANNOT be D as if you restarted the server, the daemon would automatically be reloaded. You also would not see the job executed at all if there was an issue with the daemon recognizing the changes. There is no need to manually reload the daemon as other answers have claimed.

upvoted 2 times

🗨️ 👤 **Alizadeh** 1 year, 10 months ago

Selected Answer: C

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

upvoted 1 times

🗨️ **linux_admin** 2 years, 4 months ago

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

The issue of the job only running daily is likely due to an incorrect OnCalendar schedule in the timer definition. The OnCalendar directive is used in the timer definition to specify the schedule on which the timer should run. If the schedule is incorrect, the timer will not run as expected.

For example, if the administrator wants the job to run every two hours, the OnCalendar directive in the timer definition should be set to `*:0/2`.

To resolve the issue, the administrator should check the timer definition and make sure that the OnCalendar schedule is set correctly. After making the necessary changes, the administrator should reload the timer and service definitions and check the log file again to see if the job is running as expected.

upvoted 4 times

🗨️ **linux_admin** 2 years, 4 months ago

In the previous question, the answer choice D was "The system-daemon services need to be reloaded." This answer is not the most likely cause of the issue because reloading the system-daemon services will not fix an incorrect OnCalendar schedule in the timer definition.

The OnCalendar directive is used in the timer definition to specify the schedule on which the timer should run. If the schedule is incorrect, the timer will not run as expected, regardless of the state of the system-daemon services.

The most likely cause of the issue, as stated in answer choice C, is an incorrect OnCalendar schedule in the timer definition. To resolve the issue, the administrator should check the timer definition and make sure that the OnCalendar schedule is set correctly.

upvoted 4 times

🗨️ **KnifeClown1** 2 years, 4 months ago

Selected Answer: C

When using systemd timers, the "OnCalendar" option is used to define when the job should run. By default, systemd timers are configured to use UTC time. If the administrator specified an incorrect OnCalendar schedule, the timer may not run as expected.

To schedule a job to run every two hours, the correct OnCalendar schedule should be `"0 */2 * * *"`. This will run the job every two hours, starting at the top of the hour (i.e., at 0 minutes past the hour).

If the administrator did not specify the correct OnCalendar schedule, the timer may be set to run daily or at some other incorrect interval, which would explain the behavior observed in the log file.

upvoted 4 times

🗨️ **Nvoid** 2 years, 7 months ago

Selected Answer: C

Picking C Here.

upvoted 4 times

🗨️ **SaadiaS** 2 years, 7 months ago

`*:0/2` The task will be executed every two minutes starting from the minute 0

<https://linuxconfig.org/how-to-schedule-tasks-with-systemd-timers-in-linux>

upvoted 1 times

🗨️ **SaadiaS** 2 years, 7 months ago

Selected answer C

upvoted 2 times

🗨️ **ryanzou** 2 years, 8 months ago

Selected Answer: C

I prefer C

upvoted 4 times

🗨️ **bjornborg** 2 years, 8 months ago

Selected Answer: C

OnCalendar setting is wrong. If it were D, it wouldn't run at all, but it's running every two hours

upvoted 4 times

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 tcp 6379`
- D. `semanage port -l -t http_port_tcp 6379`

Suggested Answer: B

Community vote distribution

B (100%)

 **linux_admin** 10 months, 3 weeks ago

The semanage command is used to manage SELinux policy settings on Linux systems. In this case, the `semanage port -a -t http_port_t -p tcp 6379` command is adding a new SELinux port type for the TCP protocol and port number 6379.

Here's what each option does:

-a: The -a option stands for "add", and it is used to add a new SELinux port type.

-t http_port_t: The -t option is used to specify the SELinux type for the new port, in this case, http_port_t.

-p tcp 6379: The -p option is used to specify the protocol and port number, in this case, TCP protocol and port number 6379.

With this command, the SELinux policy is updated to allow the TCP protocol to listen on port number 6379 with the http_port_t type. This is useful in cases where the application requires the use of a non-standard port number for HTTP traffic.

upvoted 3 times

 **Aamm033** 1 year, 2 months ago

Selected Answer: B

Correct answer is B -a means Add and -t Type

upvoted 3 times

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

- A. `hostnamectl status --no-ask-password`
- B. `hostnamectl set-hostname "$(perl -le "print "A" x 86)"`
- C. `hostnamectl set-hostname Comptia-WebNode -H root@192.168.2.14`
- D. `hostnamectl set-hostname Comptia-WebNode --transient`

Suggested Answer: C

Community vote distribution

C (100%)

linux_admin Highly Voted 2 years, 4 months ago

The `hostnamectl set-hostname Comptia-WebNode -H root@192.168.2.14` command sets the hostname of a Linux system.

Here's what each option does:

`set-hostname Comptia-WebNode`: The `set-hostname` option sets the hostname of the system to "Comptia-WebNode".

`-H root@192.168.2.14`: The `-H` option is used to specify the remote host to connect to. In this case, the administrator is connecting to the remote host with the IP address 192.168.2.14 as the root user.

This command sets the hostname of the remote host with IP address 192.168.2.14 to "Comptia-WebNode" as the root user. The new hostname will persist across reboots unless the administrator changes it again in the future.

upvoted 9 times

linux_admin 2 years, 4 months ago

In the context of the question, a "tag" refers to a label or identifier that is used to distinguish or identify a specific object, in this case, the web server. The tag is used to make it easier for end users to connect to the web server, as they can use the tag instead of an IP address or a hostname. The tag is typically a short and descriptive string that represents the web server, such as "Comptia-WebNode" in the example.

upvoted 7 times

yomebo Highly Voted 11 months, 2 weeks ago

Selected Answer: C

Honestly I just went for C since its the only one with an IP address..otherwise I dont really know whats going on

upvoted 5 times

Blatzzy Most Recent 1 year, 6 months ago

For adding a tag to the API, you typically don't use the `hostnamectl` command. Instead, you might want to consider adding a label or tag directly related to the API functionality. However, among the given options, the most relevant one would be:

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

This sets the transient (temporary) hostname to "Comptia-WebNode". However, keep in mind that this command is more related to the system's hostname configuration rather than specifically tagging an API.

For tagging an API, you might want to consider using metadata or labels associated with the API service itself, which is typically done in the application or service configuration rather than at the system level.

upvoted 1 times

ryanzou 2 years, 8 months ago

Selected Answer: C



To execute a `hostnamectl` command on a remote system, use the `-H, --host`. That is change a tag remotely in my view.

upvoted 3 times

TheRealManish 2 years, 8 months ago

What the heck is the question asking? Answer C, doesn't add a "tag". it assigned a hostname to a remote server. Does anyone understand this question?

upvoted 1 times

  **Nvoid** 2 years, 7 months ago

ya probably not on the test, or the question/answers will be different to actually make sense.

upvoted 1 times

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/`

Suggested Answer: C

Community vote distribution

C (67%)

D (33%)

linux_admin Highly Voted 2 years, 4 months ago

C. `rsync -a /data remote:/backup/`

The `rsync` command is used to copy files and directories from one location to another, either locally or over a network. The `-a` option is used to preserve the original file attributes, such as timestamps, permissions, and symbolic links.

In this case, the administrator is using `rsync` to copy the contents of the local directory `/data` to the remote server `remote` at the location `/backup/`. The `rsync` command will copy all the files and subdirectories in `/data` to the remote server and preserve their original attributes.

upvoted 8 times

linux_admin 2 years, 4 months ago

The other options are not correct or do not achieve the desired effect. The `scp` command is used to copy files securely over a network, but it does not preserve the original file attributes. The `ssh` command is used to log into a remote server, but it does not copy files. The `cp` command is used to copy files locally, but it does not copy files over a network.

upvoted 4 times

linux_admin 2 years, 4 months ago

Option A (`scp -p /data remote:/backup/data`) is not the correct command to achieve the desired effect of backing up the `/data` directory and its contents to the remote server.

The `scp` command is used to copy files and directories between hosts over SSH, but it does not support copying directories recursively by default. Therefore, the `/data` directory will not be backed up completely using this command.

To copy directories recursively, we would need to use the `-r` option with `scp`. In addition, we need to specify the destination directory as `/backup` instead of `/backup/data` since we want to copy the contents of the `/data` directory to `/backup`. Therefore, the correct `scp` command would be:

```
scp -r /data remote:/backup
```

upvoted 5 times

Aj26a Most Recent 1 year ago

Selected Answer: D

C is a valid option if you need to set the hostname on a remote machine. If temporary, D.

upvoted 1 times

bongobo 1 year, 3 months ago

A works, but result will be `/backup/data/data`

upvoted 1 times

Alizadeh 1 year, 10 months ago

Selected Answer: C

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

upvoted 1 times

🗨️ 👤 **Deuteronomy** 2 years, 4 months ago

The correct answer is still C. scp -p option does not copy sub-directories. Therefore, if there were sub-directories under /backup/data, it would not copy the entire contents. In other words, the option scp -r had to be used for A to be the correct answer.

upvoted 2 times

🗨️ 👤 **MissAllen** 2 years, 7 months ago

C is correct. With the -a option, rsync will copy the entire data directory with contents to the /backup directory. No need to specify the data directory.

upvoted 1 times

🗨️ 👤 **Aamm033** 2 years, 8 months ago

Selected Answer: C

C is correct...rsync is for remote copy and the tag -a is for archive mode which is good for backing up.

upvoted 1 times

🗨️ 👤 **TheRealManish** 2 years, 8 months ago

Are you sure it's C? it looks like it will back it up into the wrong remote directory. if we were using rsync we would need: srync -a /data remote:/backup/data ? wouldn't we? seems to me that option A, is the only one that would accomplish the goal?

upvoted 2 times

🗨️ 👤 **Veteran903** 2 years, 7 months ago

same thing I was thinking, definitely A

upvoted 1 times

🗨️ 👤 **Veteran903** 2 years, 7 months ago

correction!!!, Miss Allen is right

rsync: A very flexible network enabled syncing tool, a copy that doesn't send things that are already at the destination and in case of connection interruption, can pick up quickly by reissuing the same command. A great utility in terms of taking the backup or doing the incremental transfer.

scp: A raw copy, its a dump copy that absolutely copies what you ask. A very good utility for copying the files that are of small size.

upvoted 1 times

🗨️ 👤 **TheRealManish** 2 years, 7 months ago

thanks, i agree. i just tested it in my lab, and it moved the directory into the destination, so the directory integrity was kept.

upvoted 1 times

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/linux+-.git`
`git checkout -b <new-branch>`

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ **linux_admin** Highly Voted 1 year, 4 months ago

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

The `git clone` command is used to clone a remote Git repository, which in this case is the repository located at `https://github.com/comptia/linux+.git`. This command will download the entire repository to the local machine.

After cloning the repository, the administrator should create a new branch using the `git checkout -b <new-branch>` command. This will create a new branch in the Git repository and make it the current branch. The administrator can then make changes to the IaC declaration templates in this branch without affecting the main branch.

upvoted 8 times

🗳️ **linux_admin** 1 year, 4 months ago

Once the changes are made, the administrator can then commit them to the new branch using `git commit`. Finally, the administrator can merge the changes back into the main branch using `git merge`. This will allow the changes to be tracked and maintain version control of the IaC declaration templates.

The other options are not correct or do not achieve the desired effect. The `git push origin` command is used to push changes to a remote repository, but it should not be used before making changes. The `git fetch New-Branch` command is used to retrieve changes from a remote repository, but it does not create a new branch. The `git status` command is used to check the current status of the Git repository, but it does not create a new branch or allow changes to be made.

upvoted 7 times

🗳️ **Alizadeh** Most Recent 10 months, 2 weeks ago

Selected Answer: D

D. `git clone https://github.com/comptia/linux+-.git`
`git checkout -b <new-branch>`
upvoted 1 times

🗳️ **Pinnubhai** 1 year, 5 months ago

Selected Answer: D

`git checkout -b|-B <new_branch> [<start point>]`
Specifying `-b` causes a new branch to be created as if `git-branch(1)` were called and then checked out.
upvoted 1 times

🗳️ **MrGykz** 1 year, 7 months ago

You could omit `<branch>`, in which case the command degenerates to "check out the current branch", which is a glorified no-op with rather expensive side-effects to show only the tracking information, if exists, for the current branch.

`git checkout -b <new-branch> [<start-point>]`

Specifying `-b` causes a new branch to be created as if `git-branch[1]` were called and then checked out. In this case you can use the `--track` or `--no-track` options, which will be passed to `git branch`. As a convenience, `--track` without `-b` implies branch creation; see the description of `--track` below.

documentation:



<https://git-scm.com/docs/git-checkout>

upvoted 1 times

  **TheRealManish** 1 year, 7 months ago

Can anyone confirm the answer is D? I've been researching this for 3 hours and have learned the basics of GIT and have no clue if D is correct.

upvoted 1 times

  **Nvoid** 1 year, 7 months ago

I know its not A or C, so it has to be B or D, and `"-b <new_branch>"` seems right from the last time i used git in ages now.

upvoted 1 times

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

Suggested Answer: D

Community vote distribution

D (100%)

 **linux_admin** Highly Voted 1 year, 4 months ago

You can use the chmod command with the -t option to remove the sticky bit:

```
chmod -t /path/to/directory
```

You cannot remove a sticky bit from a file directly. The sticky bit is a property of a directory, not a file. When a sticky bit is set on a directory, it affects the behavior of file deletion within that directory. A file within a directory with a sticky bit set cannot be deleted by a user who does not have write permissions to the directory, regardless of the permissions on the file itself. Therefore, to remove a sticky bit, you need to use the path to the directory for which you want to remove the sticky bit, not a file within the directory.

upvoted 5 times

 **Alizadeh** Most Recent 10 months, 2 weeks ago

Selected Answer: D

The correct answer is D. chmod -t files


upvoted 1 times

 **ryan zou** 1 year, 8 months ago

Selected Answer: D

D is correct

upvoted 2 times

 **Aamm033** 1 year, 8 months ago

Selected Answer: D

D is correct the directory has Sticky bit on. in the first ls -ld you can see the "t" instead of an "x"
upvoted 2 times

Which of the following commands will display the operating system?

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

Suggested Answer: C

Community vote distribution

C (92%)

8%

 **linux_admin** Highly Voted 1 year, 4 months ago

Selected Answer: C

The correct command to display the operating system is C. `uname -o`.

The `uname` command is used to print system information. It has several options that can be used to print different types of information.

Here's what each option does in the context of the given question:

- A. `uname -n`: prints the hostname of the machine.
- B. `uname -s`: prints the kernel name of the machine.
- C. `uname -o`: prints the operating system of the machine.
- D. `uname -m`: prints the machine hardware name.

Therefore, the correct option to display the operating system is C (`uname -o`).

upvoted 5 times

 **BryanSME** Most Recent 7 months, 2 weeks ago

```
[bryan@almalinux etc]$ uname -n
almalinux
```

```
[bryan@almalinux etc]$ uname -o
GNU/Linux
```

```
[bryan@almalinux etc]$ uname -s
Linux
```

```
[bryan@almalinux etc]$ uname -m
x86_64
```

upvoted 1 times

 **examtopics11** 1 year, 1 month ago

Selected Answer: C

C. `uname -o` (OS)

-s (kernel) is Linux

-o (OS) is x86_64

<https://www.youtube.com/watch?v=1AoKTorIEDc> (2:28)

upvoted 1 times

 **Lwarder1** 1 year, 4 months ago

Selected Answer: C

`uname` man page

-a, --all

print all information, in the following order, except omit -p and -i if unknown:

-s, --kernel-name

print the kernel name

-n, --nodename

print the network node hostname

-r, --kernel-release

print the kernel release

-v, --kernel-version

print the kernel version

-m, --machine

print the machine hardware name

-p, --processor

print the processor type or "unknown"

-i, --hardware-platform

print the hardware platform or "unknown"

-o, --operating-system

print the operating system

--help

display this help and exit

--version

output version information and exit

upvoted 4 times

🗨️ 👤 **linux_admin** 1 year, 4 months ago

Selected Answer: B

The option -o for the uname command does not display the operating system. Instead, it is used to display the operating system name, which is not the same as the operating system itself. The -s option is used to display the operating system name, which typically includes the name, version, and architecture.

For example, if you run the command `uname -s` on a Linux system, it will return the following output:

Linux

So, in this case, the correct option for displaying the operating system would be -s, not -o.

upvoted 1 times

🗨️ 👤 **bjornborg** 1 year, 7 months ago

Selected Answer: C

`uname -o`

GNU/Linux

upvoted 1 times

🗨️ 👤 **Aamm033** 1 year, 8 months ago

Selected Answer: C

C is correct

upvoted 1 times

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?

- `pvcreate -L 1G /dev/app`

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

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

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

D. `mkfs.xfs /dev/app_vg/app`
`mount /dev/app_vg/app /opt/app`

Suggested Answer: D

Community vote distribution

D (72%)


B (28%)

 **Nvoid** Highly Voted 1 year, 1 month ago

Selected Answer: D

Question should be fixed to say: mounted under "/opt/app" not "/ops/app".

upvoted 10 times

 **Nvoid** 1 year, 1 month ago

Still needs to be fixed mods.

upvoted 2 times

 **linux_admin** Highly Voted 10 months, 2 weeks ago

Selected Answer: D

D. `lvcreate -L 1G -n app app_vg`

`mkfs.xfs /dev/app_vg/app`

`mount /dev/app_vg/app /opt/app`

This 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.

upvoted 8 times

 **linux_admin** Most Recent 10 months, 3 weeks ago

Selected Answer: B

Answer B is a correct set of commands to create and mount a new 1GB XFS filesystem under the directory "/opt/app".

The first command `parted /dev/sdb --script mkpart primary xfs 1GB` creates a new partition on the device "/dev/sdb" with a file system type of XFS and a size of 1GB.

The second command `mkfs.xfs /dev/sdb` creates a XFS file system on the newly created partition "/dev/sdb".

The third command `mount /dev/sdb /opt/app` mounts the newly created XFS file system to the directory "/opt/app".

This set of commands will work if `/dev/sdb` is a valid block device and there is enough disk space available to create a new partition of 1GB. It is also important to note that these changes are only temporary and will not persist after a reboot unless they are added to the `/etc/fstab` file.

upvoted 2 times

🗋️ 👤 **linux_admin** 10 months, 2 weeks ago

****Discard****

upvoted 2 times

🗋️ 👤 **bjornborg** 1 year, 1 month ago

Selected Answer: B

i get B because the others have to do with "trying" to expand a logical volume. The question says just to simply add a 1GB partition

upvoted 3 times

🗋️ 👤 **Veteran903** 1 year, 1 month ago

No, read the question again, you have to add a "file system", D is correct

upvoted 2 times

🗋️ 👤 **TheRealManish** 1 year, 1 month ago

I tested D and it worked.

upvoted 1 times

🗋️ 👤 **TheRealManish** 1 year, 1 month ago

B does not work -tested. the `mkpart` command requires an END block. if you run this command by leaving the `--script` out, you will get prompts.. I would prefer B as an answer, but it doesn't work. so D it is.

upvoted 2 times

🗋️ 👤 **[Removed]** 10 months, 3 weeks ago

can't be B, after running parted the `partprobe` command needs to be run to refresh the kernel data

upvoted 1 times

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

Suggested Answer: C

Community vote distribution

A (50%)

D (33%)

C (17%)

🗳️ 👤 **Rob74613** Highly Voted 2 years, 1 month ago

Selected Answer: D

unzip -v only displays the contents with verbosity of a compressed file but not extract

bzip2 -z will compress a file not extract

gzip will only compress a file not extract

funzip is the ONLY option here that if used just like it is will extract a file. The funzip command is used to decompress files that have been compressed with the ZIP compression algorithm. It is part of the Info-ZIP suite of tools and is typically available on Linux and Unix systems.

upvoted 9 times

🗳️ 👤 **TheMichael** Most Recent 4 months, 1 week ago

Selected Answer: A

I'm not really sure why everyone is disagreeing on this answer. It's either A or D, with A seeming to be more likely. Not sure why the top voted thinks unzip doesn't extract, maybe they are confused by the verbose flag? but it totally does extract by default.

A. Unzip -v decompresses and then extracts the files and directories contained.

the -v just shows the listing of the files as they are being extracted

B. bzip2 -z is just doing compression

C. gzip is also just doing compression here (there is no -d to decompress and the -d wouldn't extract, just decompress)

D. funzip is only designed to decompress and extract .zip files, whereas the question asks for a undisclosed format. This makes it more likely to be A.

upvoted 1 times

🗳️ 👤 **JRS99** 6 months, 2 weeks ago

Selected Answer: D

Correct answer: D

unzip -v is not sufficient to accomplish the task.

bzip2 -z is going to compress the file.

gzip is going to compress the file. (The option is gzip, not gunzip people. Don't assume.)

funzip alone will extract compressed ZIP files.

upvoted 2 times

🗳️ 👤 **insanegrizly** 8 months ago

Selected Answer: D

unzip -v: The -v option with unzip is used to list the contents of a zip file but does not extract them. So, while unzip is a valid command for extraction, -v alone won't do it, therefor i'm going with D.

upvoted 1 times

🗳️ 👤 **Qubert2** 8 months, 3 weeks ago

funzip is a command in Linux for decompression. It is part of the Info-ZIP suite and is used to extract the contents of a single file from a compressed .zip archive, sending the decompressed data to standard output (stdout). Unlike the unzip command, funzip works only with a single member from a zip archive or from a stream of compressed data, making it useful in pipelines or when you want to process the decompressed content on-the-fly without fully extracting the file to disk.

upvoted 1 times

🗳️ 👤 **riddie78** 1 year, 5 months ago

MOST LIKELY GUNZIP...

the g and f are close on keyboard... misspelling most likely... so D

upvoted 3 times

🗨️ **kennethcan** 1 year, 8 months ago

Selected Answer: D

As written, "funzip" is the only command here that will actually extract a .zip. None of these will extract anything else compressed without additional options. Hopefully this question is just wrong.

upvoted 4 times

🗨️ **linux_admin** 2 years, 4 months ago

Selected Answer: C

C. gzip is the best option.

upvoted 3 times

🗨️ **linux_admin** 2 years, 4 months ago

Selected Answer: A

A. unzip -v

upvoted 1 times

🗨️ **linux_admin** 2 years, 4 months ago

Discard

upvoted 1 times

🗨️ **bjornborg** 2 years, 7 months ago

Selected Answer: A

opposite of gzip is gunzip...gzip compresses, gunzip uncompresses ... just because .gz suffix on a file, it will not uncompress with gzip...needs gunzip

gzip file.gz

gzip: file.gz already has .gz suffix -- unchanged

upvoted 4 times

🗨️ **TheRealManish** 2 years, 7 months ago

gzip -d decompresses,, but then they aren't giving us the -d in the answer.. These questions are ridiculous. we shouldn't have to be guessing if an option can or can't be used.

upvoted 2 times

🗨️ **Veteran903** 2 years, 7 months ago

gzip is a file format and a software application used for BOTH, file compression and decompression.

upvoted 1 times

🗨️ **MissAllen** 2 years, 7 months ago

I agree with A. Be careful in assuming that the package being downloaded is in gzip format. It can be in any acceptable format in Linux, such as xz or bz2 or zip. Also, the wording of the question is tricky. When I took an earlier version of Linux+ (when it was two exams), they were pretty clear when they wanted just a command as an answer versus the entire command syntax, which can include options and arguments.

upvoted 2 times

🗨️ **Aamm033** 2 years, 8 months ago

Selected Answer: A

Answer is A because unzip -v decompress with verbose and gzip without the -d option will compress.

upvoted 4 times

🗨️ **TheRealManish** 2 years, 8 months ago

I disagree and think the answer is C. the question says it download a software package. a linux software package is going to be stored in gz format not windows zip format. The question also doesn't ask for an argument/operation, just which command.

upvoted 1 times

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 kernel scope link src 5.189.153.89
```

Server output 1:

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

Server output 2:

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

Server output 3:

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

Server output 4:

```
default via 104.21.75.254 dev eth0
104.21.75.0/24 dev eth0 proto kernel 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.

Suggested Answer: B

Community vote distribution

🗳️ 👤 **Nvoid** Highly Voted 2 years, 7 months ago

Selected Answer: C

It's a Firewall issue, i'm picking `C`.

Please note, it's easy to misread the title of the outputs, read carefully -i made the same mistake from `TheRealManish` contribution also...

it reads:

Server Output#1

Server Output#2

Server Output#3

Server Output#4

NOT

Server #1 Output

Server #2 Output

Server #3 Output

Server #4 Output

Cheers!

upvoted 7 times

🗳️ 👤 **TheRealManish** 2 years, 7 months ago

OMG thanks so much! i totally missed the wording!! i was like why the F is it telling us all of these other outputs!

upvoted 2 times

🗳️ 👤 **TheRealManish** 2 years, 7 months ago

Thanks again, but im reading thru this and the firewall rejection is to reject with ICMP.. but instead it is rejecting with TCP reset. It seems like C is also wrong, but its way closer than all of the rest.

upvoted 2 times

🗳️ 👤 **Nvoid** 2 years, 7 months ago

TCP/IP defines the kernel will send an ICMP message back with an "Port unreachable" message for UDP services, and TCP RST messages for TCP - REF: <https://unix.stackexchange.com/questions/261360/icmp-port-unreachable-error-even-if-port-is-open>

So from what i can tell, TCP response from the server should be sending a RST not a ICMP "Port unreachable" message, thats for udp!

upvoted 1 times

🗳️ 👤 **Nvoid** 2 years, 7 months ago

Glad we're working together, i figured you would have taken your test already and passed.

upvoted 1 times

🗳️ 👤 **Aj26a** Most Recent 1 year ago

Selected Answer: C

From the provided outputs and error message, the connectivity issue seems to be caused by the server's firewall configuration.

Specifically, Server Output 1 shows that the firewall on the server is configured to reject SSH connections (port 22) from several IP addresses, including 5.189.153.89, which matches the IP address of the workstation.

upvoted 3 times

🗳️ 👤 **BryanSME** 1 year, 7 months ago

This statement on Server output 2: "Active: active (running) since" means that even tho the service did not load at startup, it has been manually started,

manually starting `sshd`:

➤ `sudo systemctl start sshd`

➤ `sudo systemctl status sshd`

• `sshd.service` - OpenSSH Daemon

Loaded: loaded (/usr/lib/systemd/system/sshd.service; enabled; vendor preset: disabled)

Active: active (running) since Tue 2020-07-14 11:13:08 -03; 2s ago

From <https://bbs.archlinux.org/viewtopic.php?id=257365>

That leaves the firewall issue as correct.

upvoted 1 times

🗳️ 👤 **Cieliog** 1 year, 8 months ago

Selected Answer: B

Comptia is famous for these gotcha questions. C is probably correct, but B is probably the Comptia answer.

upvoted 2 times

🗳️ 👤 **clair12** 8 months, 2 weeks ago

might not be B actually, sshd seems to have been running for the past two weeks

upvoted 1 times

🗳️ 👤 **mrtwister76** 2 years, 1 month ago

Firewall issue

upvoted 1 times

🗳️ 👤 **Rob74613** 2 years, 1 month ago

Selected Answer: C

As the question states the user tried to ssh from ONE of the workstations, and one of the workstations (Workstation 1) has an IP that is being rejected in the firewall with port 22 (aka ssh)

upvoted 1 times

🗳️ 👤 **ominousred** 2 years, 2 months ago

I selected "B" because if you take any service and DISABLE it, it will not work. In this case, ssh is disabled.

upvoted 1 times

🗳️ 👤 **CodeMaestro** 2 years, 1 month ago

Yes the ssh service is disabled but as you can see it has been running, what that means is that the service does not start up automatically but needs the administrator to actively start it up, but the firewall drops the packets from port 22 and thus it is a firewall issue.

upvoted 2 times

🗳️ 👤 **linux_admin** 2 years, 4 months ago

Selected Answer: C

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

upvoted 1 times

🗳️ 👤 **bjornborg** 2 years, 7 months ago

Selected Answer: C

Server output 2 -> sshd "active (running)", so answer not B :-(

Server output 1 -> port 22 being blocked from all sources ... firewall issue

Everything else looks fine

upvoted 3 times

🗳️ 👤 **TheRealManish** 2 years, 7 months ago

MODS, I think i might have hit the flag button by accident, please disregard .

the output says we are connecting to the server ending in .76. Thats server 3. so we ignore all of the output for server 1,2,4. All we know about 3 is that it has an IP address. If a service is not running on a machine, it will send a connection refusal..

upvoted 1 times

🗳️ 👤 **TheRealManish** 2 years, 8 months ago

It seems like B to me.. as Server 3 seems to set the correct IP. receiving a reset packet back indicates a service is not running.

upvoted 1 times

🗳️ 👤 **Veteran903** 2 years, 7 months ago

server output is telling you ssh is active and running, also, as you can see the firewall is rejecting all connections so this is clearly a firewall issue, definitely C

upvoted 1 times

🗳️ 👤 **TheRealManish** 2 years, 7 months ago

are you sure? it says we are connecting to host 104.21.75.76. thats server 3. therefore the ONLY output we should concern ourselves with is the output from server 3. the server 3 output is super vague. no talk about firewall or ssh port. i wish we had a way to get together and review this stuff :)

upvoted 1 times

🗳️ 👤 **TheRealManish** 2 years, 7 months ago

disregard my comment here, apparently, I can't read lol.

upvoted 1 times

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

Suggested Answer: A

Community vote distribution

A (100%)

 **bjornborg** Highly Voted 2 years, 7 months ago

Selected Answer: A

amazingly, it's the same file on both Debian and Red Hat class systems
upvoted 5 times

 **Qubert2** Most Recent 8 months, 3 weeks ago

A:

upvoted 1 times

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

Suggested Answer: D

Community vote distribution

D (83%)

A (17%)

🗳️ 👤 **Aj26a** 1 year ago

Selected Answer: D

Given the key features mentioned – portability, high availability, and scalability in production – the best choice for the new design is:

D. Kubernetes

Kubernetes is an open-source platform designed for automating deployment, scaling, and operations of application containers across clusters of hosts. It provides:

Portability: Kubernetes can run on various environments, including on-premises, public cloud, and hybrid cloud.

High Availability: Kubernetes supports self-healing capabilities, automatic restarts, and load balancing.

Scalability: Kubernetes can scale applications automatically based on demand.

While Docker provides containerization, Kubernetes is specifically designed to manage containerized applications at scale and across multiple hosts, making it the best choice for the specified requirements.

So, the correct answer is D. Kubernetes.

upvoted 3 times

🗳️ 👤 **Monty97** 1 year, 3 months ago

Selected Answer: D

Kubernetes is a container orchestration platform that provides features such as automatic scaling, load balancing, and self-healing capabilities. It allows for the management of containerized applications across multiple hosts, making it suitable for environments requiring high availability and scalability.

Docker is a containerization platform that allows you to package and distribute applications and their dependencies in containers. While Docker is commonly used for containerization, it lacks built-in orchestration features like those provided by Kubernetes.

On-premises systems and cloud-based systems refer to infrastructure deployment models rather than containerization technologies. They can be used in conjunction with container orchestration platforms like Kubernetes.

upvoted 3 times

🗳️ 👤 **imnewtothis** 1 year, 3 months ago

Selected Answer: D

Kubernetes is a powerful container orchestration platform that provides features such as portability, high availability, and scalability in production environments. It automates the deployment, scaling, and management of containerized applications, allowing for efficient resource utilization and seamless scaling. Kubernetes can run on both on-premises systems and cloud-based systems, providing flexibility in deployment options. While Docker is a popular containerization platform, Kubernetes extends its capabilities by providing advanced orchestration features suitable for production environments. On-premises and cloud-based systems are infrastructure choices, whereas Kubernetes is a platform for managing containerized applications regardless of the underlying infrastructure.

upvoted 1 times

🗳️ 👤 **DRVision** 1 year, 6 months ago

Selected Answer: A



"creating this type of deployment", kubernetes is used for container orchestration

upvoted 1 times

  **DRVision** 1 year, 6 months ago

kubernetes can be used to enhance docker or cloud based services but cannot be used as a stand alone deployment, and cloud based services are not as portable as containers, i.e.docker

upvoted 1 times

  **nixonbii** 2 years, 4 months ago

Do not understand the context of this question. Both MS Azure and AWS offer Kubernetes based container services. It seems that the cloud is far more portable, accessible, and available than an on-premises system which would rely on VPN tunnels over public internet to access the resources. What are the assumptions we are supposed to make when posed with a question like this? Also, why not Docker?

upvoted 2 times

  **linux_admin** 2 years, 4 months ago

For the given scenario where the administrator is looking for portability, high availability, and scalability in production, the best choice would be to use Kubernetes. Kubernetes is an open-source platform that provides these features and helps to automate the deployment, scaling, and management of containerized applications. It is widely used in production environments and provides a lot of flexibility in terms of managing resources and scaling the application.

upvoted 1 times

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

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

Suggested Answer: C

Community vote distribution

C (100%)

  **linux_admin** Highly Voted 10 months, 3 weeks ago

Selected Answer: C

C. Jenkins is commonly used for creating CI/CD (Continuous Integration/Continuous Deployment) pipelines. It is an open-source automation server that can be used to automate various tasks, including building, testing, and deploying software applications. Other tools like Chef, Puppet, and Ansible are typically used for configuration management and infrastructure automation.

upvoted 9 times

  **LouSassle** 8 months ago

Ansible is an orchestration tool that is commonly used for Red Hat Enterprise Linux deployments, is agentless, and relies on the Python programming language.

Chef is an orchestration tool that utilizes agents and the Ruby programming language.

Puppet is an orchestration tool that is agentless and recognizes configuration files in Bash, Python, Ruby, YAML, and PowerShell.

upvoted 2 times

  **ryanzou** Most Recent 1 year, 2 months ago

Selected Answer: C

C is correct, CI/CD build/Testing step.

upvoted 2 times

  **Aamm033** 1 year, 2 months ago

Selected Answer: C

C is correct... other options are for config management.

upvoted 2 times

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`

Suggested Answer: B

Community vote distribution

C (95%)

5%

linux_admin Highly Voted 2 years, 4 months ago

Selected Answer: C

The correct answer is C. The "echo "umask 377" >> /home/web/.bashrc" command will satisfy the requirement by adding the umask setting to the user's .bashrc file. The umask setting determines the default file permissions for newly created files and directories. The value of 377 sets the default permissions to read-only for the owner of the file. The umask is specified in octal notation, and each digit represents the permissions for a different category of users. In this case, the first digit represents the permissions for the owner of the file, the second digit represents the permissions for the group owner of the file, and the third digit represents the permissions for others. A value of 3 means that the owner has read and write permissions, and a value of 7 means that the owner has read, write, and execute permissions. By setting the umask to 377, the owner of newly created files will have read-only permissions, the group owner will have no permissions, and others will have no permissions.

upvoted 5 times

clair12 8 months, 2 weeks ago

why do the group and others not have permissions?

doesnt 7 mean rwx?

upvoted 2 times

bc1235813 Most Recent 1 year, 3 months ago

Selected Answer: C

The question states files that "are" created not "were" created. "B" takes care of files that "were" created/exist, "C" takes care of files that "are" (implying in the future) created. Semantics of the english language.

upvoted 2 times

e418137 1 year, 4 months ago

Selected Answer: C

Demonstrable: `UMASK=$(umask); umask 377; touch file; mkdir dir; umask $UMASK; unset UMASK; ls -l``

upvoted 1 times

DRVision 1 year, 6 months ago

Selected Answer: C

B is incorrect as the -R makes it recursive and it will change all the contents of already created files but will not affect any new files

C is correct the correct option

upvoted 1 times

DRVision 1 year, 6 months ago

Here's how you can convert the octal number 377 to binary:

3 in octal is 011 in binary.

7 in octal is 111 in binary.

7 in octal is 111 in binary.

011 means read (r) permission is denied and write (w) and execute (x) permissions are granted.

111 means read (r), write (w), and execute (x) permissions are granted.

111 means read (r), write (w), and execute (x) permissions are granted.

upvoted 3 times

🗨️ **DRVision** 1 year, 6 months ago

Therefore, a umask of 377 typically results in new files having read-only permissions for the owner, and no permissions for the group and others.

and b appending this umask to the user web bashrc file, anything they create will have a default read permissions. Echo works like a print statement, not displays what the user wants...lol

upvoted 1 times

🗨️ **cruxty** 1 year, 9 months ago

Echo doesn't change anything it just displays what the user wants in the echo requests output

Echo Answer is B answer is b

upvoted 3 times

🗨️ **Rob74613** 2 years, 1 month ago

Selected Answer: B

The permission value of 400 translates to read-only permission for the owner and no permission for the group and others. By using this command, all files created by the user web within the /home/web directory will have read-only permission for the owner, which meets the requirement. However its highly voted that "C" is the answer but I believe its wrong for this reason.

The "umask" command is used to set the default permissions for new files and directories created by the user. In this case, the value "377" is being set as the umask, which means that any new files or directories created by the user will have permissions set to 600 (read and write access for the owner, and no access for group members or others). Which question specifically mentions "Read-Only"

upvoted 1 times

🗨️ **Rob74613** 2 years, 1 month ago

Disregard, I'm an idiot, dont listen to me...I change my answer to C

upvoted 1 times

🗨️ **luken7777** 2 years ago

why?!!! you were absolutely right with your first answer. Anyone who think that the right answer is C should think really hard about their life.

"umask 377" - in binary

3 = 011

7 = 111

7 = 111

Then, we invert each bit to get the bits that should be turned off:

011 -> 100

111 -> 000

111 -> 000

And there you have 100 which in decimal is 600 (4 - read, 2 write).

Question says that the owner of the file should have read permission only - 400, so that automatically excludes C as the answer.

upvoted 5 times

🗨️ **luken7777** 2 years ago

Update: And there you have 100 000 000 which in decimal is 600 (4 - read, 2 write).

upvoted 2 times

🗨️ **Rob74613** 2 years ago

Incorrect, you're off a tad here, since we're working with a directory they umask is set to 777

777 -> 111 111 111

377 -> 011 111 111

100 000 000

When you convert 100 from binary to decimal you get 4 (not 6). Your calculating the binary left to right instead of right to left. 00000100 = 2^2

Also: <https://www.linuxtrainingacademy.com/all-umasks/>

upvoted 2 times

🗨️ 👤 **angellorv** 2 years, 3 months ago

Answer B: chmod changes the permissions of each given file according to mode; -R (recursive flag - successive executions); 400 = owner "read", group "---", others "---"; to the home directory of user "web"

Every file created by "web" will have 400 permissions after this point

upvoted 3 times

🗨️ 👤 **nixonbii** 2 years, 4 months ago

I found the umask field on CentOS8 in the /etc/profile file, nothing in .bashrc. On Ubuntu, no umask field in .bashrc, /etc/profile or /etc/bash.bashrc. I guess it has to be inserted if you want to apply it.

upvoted 1 times

🗨️ 👤 **nixonbii** 2 years, 4 months ago

Looked at the .bashrc file and reviewed the study guide. The file is used to configure bash shell for a given user. Found no stubs or commented out sections for the umask command. Can someone explain?

upvoted 1 times

🗨️ 👤 **Pinnubhai** 2 years, 5 months ago

Selected Answer: C

answer is C

upvoted 2 times

🗨️ 👤 **poni0331** 2 years, 6 months ago

Selected Answer: C

C is correct here

upvoted 2 times

🗨️ 👤 **Nvoid** 2 years, 7 months ago

Selected Answer: C

Charlie mods!!

upvoted 2 times

🗨️ 👤 **Veteran903** 2 years, 7 months ago

Tested on centOS 9, answer is C

upvoted 3 times

🗨️ 👤 **bjornborg** 2 years, 7 months ago

Selected Answer: C

yep, C

upvoted 2 times

🗨️ 👤 **MissAllen** 2 years, 7 months ago

I agree with C, however I would have preferred the umask value be set at 0277. The default mode for directories is 777, so applying umask 0377 leaves you with octal 400 for directory creation. The question is asking about creating new files, so applying 0377 to the default mode for files of 666, you get 300 for the permissions, which is -wx for the web user. Not quite read.

upvoted 4 times

🗨️ 👤 **ryanzou** 2 years, 7 months ago

Selected Answer: C

C seems correct.

upvoted 1 times

🗨️ 👤 **TheRealManish** 2 years, 8 months ago

Selected Answer: C

It seems like the answer is C? the key here is that it is "all files that are created", so it is not a one time permission change.

upvoted 1 times

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_logindf.so`
- D. `pam_nologin.so`

Suggested Answer: *D*

🗉 👤 **linux_admin** 10 months, 3 weeks ago

The correct PAM module to accomplish the task of preventing logins from accounts other than root when the file `/etc/nologin` exists is `pam_nologin.so`. This PAM module checks for the existence of the `/etc/nologin` file and denies logins if the file exists. By using `pam_nologin.so` in the appropriate PAM configuration file, the systems administrator can ensure that only root logins are allowed when the `/etc/nologin` file is present.


upvoted 4 times

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

Suggested Answer: C

Community vote distribution



linux_admin **Highly Voted** 1 year, 4 months ago

Selected Answer: C

Answer B (systemctl disable nginx) and answer C (systemctl mask nginx) both achieve the goal of disabling the nginx service, but they do it in different ways.

The systemctl disable nginx command disables the nginx service so that it will not start automatically when the system boots up. However, the service can still be manually started if needed.

The systemctl mask nginx command masks the nginx service, which means that it cannot be started, either automatically or manually. This is a more secure option compared to systemctl disable nginx, as it completely prevents the service from being started.

In this scenario, either answer B or answer C would be correct, depending on the level of security that the systems administrator wants to enforce. If the administrator wants to prevent the nginx service from starting in any circumstance, then answer C (systemctl mask nginx) is the better choice. If the administrator wants to prevent the nginx service from starting automatically but still allow it to be manually started if needed, then answer B (systemctl disable nginx) is the better choice.

upvoted 6 times

BryanSME **Most Recent** 7 months, 2 weeks ago

<https://askubuntu.com/questions/816285/what-is-the-difference-between-systemctl-mask-and-systemctl-disable>

Disabling the service deletes the symlink, so the unit file itself is not affected, but the service is not loaded at the next boot, when systemd reads /etc/systemd/system.

However, a disabled service can be loaded, and will be started if a service that depends on it is started; enable and disable only configure auto-start behaviour for units, and the state is easily overridden.

A masked service is one whose unit file is a symlink to /dev/null. This makes it "impossible" to load the service, even if it is required by another, enabled service.

When you mask a service, a symlink is created from /etc/systemd/system to /dev/null, leaving the original unit file elsewhere untouched. When you unmask a service the symlink is deleted.

upvoted 1 times

Alizadeh 10 months, 2 weeks ago

Selected Answer: C

The correct answer is C. systemctl mask nginx.

upvoted 1 times

BreakOff874 1 year, 3 months ago

Selected Answer: C

systemctl mask will prevent the service from being automatically started by other services.

upvoted 2 times

KnifeClown1 1 year, 4 months ago

Selected Answer: B

The correct command to disable the nginx service from being automatically and manually started is "systemctl disable nginx". The command "systemctl disable" will disable the service, which means that it will not start automatically during boot and will not be able to be started manually.

The "systemctl stop nginx" command will only stop the service if it is currently running, but will not prevent it from being started again in the future.

The "systemctl cancel nginx" and "systemctl mask nginx" commands are not commonly used and will not accomplish the task of disabling the nginx service.

upvoted 2 times

  **KnifeClown1** 1 year, 4 months ago

Correct answer = C

upvoted 2 times

  **Notnotataco** 1 year, 7 months ago

Should the answer be B? Based on my reading, it looks like B would be the best choice.

upvoted 2 times

  **Veteran903** 1 year, 7 months ago

hello,

please, read the question again, you may not be understanding what they are asking, CompTIA wording is brutal so be careful. You can manually start a disabled service with the systemctl start command after the system boots, to prevent this, you must use the mask subcommand, masking the service links its configuration to /dev/null.

A user or process will not be able to start this service at all (whereas with a disabled service, a user or process can still start it). Use the unmask subcommand to reverse the setting:

Ex: \$ sudo systemctl mask sshd

upvoted 3 times

  **Notnotataco** 1 year, 7 months ago

I'm an idiot lol....thank you all!!!!

upvoted 1 times

  **TheRealManish** 1 year, 7 months ago

enable/disable just says what should happen at boot. it will not stop a user or service from manually executing like mask will

upvoted 1 times

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.service.
- 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.

Suggested Answer: B

 **linux_admin**  10 months, 3 weeks ago

The TimeoutStartUsec configuration determines the amount of time that the system service manager should wait for the service to start before timing out. If the logsearch.service is taking longer than the specified timeout to start, the system service manager may consider the service to have failed to start and may terminate it. Increasing the TimeoutStartUsec configuration for the logsearch.service would give it more time to start and prevent it from being terminated prematurely. This may be necessary if the service requires additional time to start due to resource constraints or other factors.

upvoted 9 times

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

Suggested Answer: D

Community vote distribution

D (100%)

Alizadeh Highly Voted 10 months, 2 weeks ago

Selected Answer: D

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.

upvoted 5 times

Damon54 Most Recent 11 months, 2 weeks ago

any comments ? please

upvoted 1 times

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 -l /dev/sdb1`
- C. `lsblk /dev/sdb1`
- D. `du -sh /oracle1`

Suggested Answer: A

Community vote distribution

A (100%)

 **linux_admin** Highly Voted 10 months, 3 weeks ago

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 `Use%` 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.

upvoted 10 times

 **Aamm033** Most Recent 1 year, 2 months ago

Selected Answer: A

Correct. flag -i for inodes instead of block

upvoted 1 times

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`

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **Alizadeh** 1 year, 10 months ago

Selected Answer: A

The correct answer is A. `dnf remove packagename`.

upvoted 1 times

🗳️ 👤 **linux_admin** 2 years, 4 months ago

If you are using a Red Hat-based Linux distribution, such as Fedora or CentOS, the correct command to remove an RPM package that was installed is `dnf remove packagename` or `rpm -e packagename`.

In this case, the command `dnf remove packagename` should be used as it is the modern package manager for Fedora and CentOS. The `rpm -e packagename` command is also correct, but `dnf` is recommended as it provides improved functionality over `rpm`.

`apt-get remove packagename` and `apt remove packagename` are commands used to remove packages on Debian-based distributions such as Ubuntu, and are not relevant for Red Hat-based distributions.

upvoted 4 times

🗳️ 👤 **EngAbood** 10 months, 2 weeks ago

but the question didn't mention the distributions :)

upvoted 2 times

🗳️ 👤 **clair12** 8 months, 2 weeks ago

It said RPM packages which are usually used on RHEL based systems. (`rpm` is redhat package manager)

upvoted 2 times

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

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **LinusSusTips** 9 months, 3 weeks ago

-l is an invalid option!

upvoted 1 times

🗳️ 👤 **linux_admin** 10 months, 3 weeks ago

The correct command to look at the last 20 lines of a log is tail -n 20.

The tail command is used to display the last part of a file, and the -n option is used to specify the number of lines to display. In this case, the administrator wants to display the last 20 lines of the log, so the correct command is tail -n 20.

upvoted 4 times

🗳️ 👤 **poni0331** 1 year ago

Selected Answer: B

B is the correct answer. Tail displays the last 10 lines by the default. The flag -n refers to the "number of lines"

upvoted 2 times

🗳️ 👤 **poni0331** 1 year ago

... followed by the number of lines specified

upvoted 1 times

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.

Suggested Answer: B

Community vote distribution

A (50%)

B (50%)

🗳️ 👤 **27c0a50** 7 months, 1 week ago

Selected Answer: B

Disk throughput is 44Mbps. sdb1 has a combined throughput of 44305Kbps (44Mbps). The answer is B.

upvoted 2 times

🗳️ 👤 **Qubert2** 8 months, 2 weeks ago

the TPS of sdb and sdb1 add up to 690, which is the disk's maximum IOPS. Therefore, the answer is A. Answer B says "permitted" - that's not a technical limitation, that's configuration.

upvoted 1 times

🗳️ 👤 **clair12** 8 months, 2 weeks ago

you cant see IOPS using this command, so those are irrelevant.

high iowait means that the system is stuck or that it's not able to provide enough resources

upvoted 1 times

🗳️ 👤 **Kashim** 10 months, 3 weeks ago

Selected Answer: B

B correct

upvoted 1 times

🗳️ 👤 **sademik** 1 year, 5 months ago

Selected Answer: B

Answer is B

upvoted 1 times

🗳️ 👤 **Damon54** 1 year, 9 months ago

Selected Answer: B

is more accurate

upvoted 1 times

  **Damon54** 1 year, 9 months ago

Selected Answer: A

Pla any comments ? A is not correct ? IOPS Input/Output Operationd per seconds ...

upvoted 3 times

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

Suggested Answer: B

Community vote distribution


B (50%)

D (50%)

 **Lwarder1** Highly Voted 2 years, 5 months ago

Can someone please explain why the answer is ip route get 192.163.1.40 from 10.0.2.15 and not either A or D? Where did 192.163.1.40 even come from?

upvoted 5 times

 **NastyNutsu** 4 months, 1 week ago

- A. this syntax is incorrect
- B. this is correct syntax, but 192.163.1.40 should be 192.168.1.10
- C. incorrect syntax
- D. this is also incorrect syntax

ip route get 192.168.1.40 from 10.0.2.15 is the answer assuming 163 is a typo (should be ...168...).

upvoted 1 times

 **JRS99** Most Recent 5 months, 3 weeks ago

Selected Answer: B

Answer B is the most up to date option.

upvoted 1 times

 **clair12** 8 months, 2 weeks ago

Selected Answer: D

as mikefnt said B has the wrong IP

upvoted 1 times

 **MikeFNT123** 9 months, 2 weeks ago

Answer is D

B displays the wrong IP. I didn't see it at first either...

upvoted 2 times

 **e418137** 1 year, 4 months ago

Selected Answer: B

"Test the route" is generic. Both B & D are correct, but `ip` supplants `route`. IOW, the `net-tools` suite (containing `route`) is becoming or has become deprecated by the `iproute2` suite (containing `ip`) in mainstream Linux distributions.

upvoted 2 times

 **sademik** 1 year, 5 months ago

Selected Answer: D

Answer is D

upvoted 1 times

 **mutawakil** 1 year, 6 months ago

A. route -e get to 192.168.1.40 from 10.0.2.15

Here's why:



route -e: This flag displays the effective routing table, including both static and dynamic routes currently in use.

get to: This specifies the type of information to retrieve, in this case, the route to the destination IP address.

192.168.1.40: This is the destination IP address for which you want to test the route.

from 10.0.2.15: This specifies the source IP address from which the route test will be initiated.

upvoted 2 times

  **angellorv** 2 years, 2 months ago

command format:

ip route get [destination ip address] from [origin ip address]

The get argument is equivalent to sending a packet along this path.

Answer B is using the correct command and I think 192.163.1.40 is a simple mistake

upvoted 2 times

  **POGActual** 2 years, 3 months ago

I dont know if the wrong IP Address was included on purpose, but neither B or C can be correct in the current format. The incorrect IP is present in the answer.

upvoted 3 times

  **linux_admin** 2 years, 4 months ago

B. ip route get 192.168.1.40 from 10.0.2.15

The "ip route get" command is used to display the path that packets take to reach a specific network host, in this case, 192.168.1.40, from the source IP address 10.0.2.15. This command will show the routing table entries used to determine the path of the packets.

upvoted 2 times

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`

Suggested Answer: B

Community vote distribution

D (100%)

linux_admin Highly Voted 1 year, 10 months ago

Selected Answer: D

D. `cat toDelete.txt | xargs rm -rf`

The "`cat toDelete.txt`" command will display the contents of the "`toDelete.txt`" file, which contains the names of the files and directories to be deleted. The "`xargs rm -rf`" command will take the standard input from the "`cat`" command and pass it as arguments to the "`rm`" command. The "`-rf`" options of the "`rm`" command are used to forcibly delete the specified files and directories, including any files and directories within them, even if they are write-protected. The "`xargs`" command is used here to process the input one line at a time and execute the "`rm`" command for each line.

upvoted 8 times

bc1235813 Most Recent 9 months, 3 weeks ago

Selected Answer: D

The `rm` command will do nothing to act on the contents of the file without help. As written `rm` will remove the file in question. You need to expose the contents of the file to `rm` in order for it to act on them. Either `cat` the file and hand it off to `xargs` or "`xargs rm -rf < file`" or `cat` it, send it to `awk` and let `awk` pass the contents to `rm`.

upvoted 1 times

e418137 10 months, 3 weeks ago

Selected Answer: D

D. ``xargs`` builds and executes command lines from standard input. Assuming the file, `toDelete.txt`, contains one file name per line, this works. The other answers are nonsense.

upvoted 1 times

Nvoid 2 years, 1 month ago

Selected Answer: D

mod correct, the question states "`sobelete.txt`" and should say "`toDelete.txt`".

upvoted 3 times

Nvoid 2 years, 1 month ago

mod correct, the answer id D.

upvoted 2 times

ryanzou 2 years, 1 month ago

Selected Answer: D

Definitely is D

upvoted 1 times

TheRealManish 2 years, 2 months ago

Selected Answer: D

Tested, confirmed D is correct

upvoted 3 times

Aamm033 2 years, 2 months ago

Selected Answer: D

D is the correct answer. The selected answer delete the file. with answer D it removes the files listed inside the file. xargs - build and execute command lines from standard input.

upvoted 3 times

  **bjornborg** 2 years, 2 months ago



I'm pretty sure this one is D. Option B just removes the file, but the question asks to remove files & directories "listed inside" the file. Just tested this out on my Linux machine ... It's D

upvoted 1 times

  **Veteran903** 2 years, 1 month ago

I didn't know one can have a directory inside a file, I thought it was the other way around, directories can contains file but file cannot contain directories, actually as B option states if you provide the options listed you can accomplish the task with no problem, i just tested it and it works fine.

upvoted 2 times

  **Nvoid** 2 years, 1 month ago

i think your getting muxed up, the "directory" is just a location inside the file "toDelete.txt".

upvoted 2 times

  **Veteran903** 2 years ago

No, a directory is always a directory, i tested B with all the options and it works perfect on CentOS 9

upvoted 1 times

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  appl

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


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

Suggested Answer: A

Community vote distribution


D (100%)

 **linux_admin** Highly Voted 1 year, 10 months ago

Selected Answer: D

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.
upvoted 7 times

 **bc1235813** Most Recent 9 months, 3 weeks ago

I agree "A" gives the process a higher priority therefore allowing it to consume even more system resources. Apparently the test developers don't listen to our comments or QC them. Just think a lot more experienced unix'ers would ace these tests if the answers were correct.
upvoted 3 times

 **mrtwister76** 1 year, 7 months ago

Definitely not A.

It's D.

upvoted 1 times


 **HMAC** 1 year, 8 months ago

it's A the question said permanently killing the process solves the immediate problem but after a reboot the process will start again
upvoted 1 times

 **tutita** 1 year, 7 months ago

when you renice to -20 you are actually giving higher priority to the process, the more -higher the number the more priority

upvoted 2 times

 **Nvoid** 2 years, 1 month ago

Selected Answer: D

I picked D here. Won't the users script will still consume those resources no matter what priority you set? even if the renice command was correct?
upvoted 1 times

 **MissAllen** 2 years, 1 month ago

Actually, I am going with D. The renice command, when using a negative number (max of -20), actually calls for the process to get more cpu cycles and therefore put more demand on the CPU. For A to be correct, the renice command would need to use a positive number which will slow it down.

Note that the output of the top command is wrong also. When you have a nice value of -10, then the priority value would be positive 10 (ten subtracted from the default of 20).

upvoted 2 times

  **TheRealManish** 2 years, 2 months ago

terrible wording, where as the renice would make the process use less CPU.. the kill -9 would also "permanently resolve the issue"

upvoted 1 times

  **Veteran903** 2 years, 1 month ago

Comptia wording is brutal, sometime I wonder if they are really testing our knowledge or just trolling thre hell out of everyone.....lol, maybe both

upvoted 5 times

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`

Suggested Answer: A

 **linux_admin** Highly Voted 10 months, 3 weeks ago

A. `chmod 4744 dev_team.txt`

The "chmod" command is used to change the access rights of a file in Linux. The "4744" options specify the access rights in binary form. The first number (4) represents the setuid (SUID) bit, which causes the file to run with the effective user ID of the owner of the file, rather than with the effective user ID of the user who is running it. The other three numbers (744) represent the permissions for the owner, group, and others, respectively. In this case, the owner has read, write, and execute permissions (7), the group has read and execute permissions (4), and others have only read permission (4).

So, the command "`chmod 4744 dev_team.txt`" sets the SUID bit and the access rights to 744 for the file `dev_team.txt`. The SUID bit allows the file to be executed with the owner's permissions, which can be useful for setuid scripts or other special purpose files.

upvoted 8 times

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/`

Suggested Answer: C

Community vote distribution

C (100%)

 **bjornborg** Highly Voted 2 years, 1 month ago

Selected Answer: C

"operation not permitted" error (usually) means the immutable attribute is set.

`chattr -i` will remove the immutable attribute

upvoted 5 times

 **ericsrz** Most Recent 1 year ago

The command `chattr -R -i data/` is used to change the attributes of a file or directory in Linux. Here's what each part of the command does:

`chattr`: This is the command used to change file or directory attributes.

`-R`: This option makes the command recursive, meaning it will apply to the specified directory and its contents.

`-i`: This option removes the 'immutable' attribute. When a file is immutable, it cannot be modified, deleted, or renamed, and no link can be created to it.

`data/`: This is the directory that the command will be applied to.

So, `chattr -R -i data/` will remove the 'immutable' attribute from the 'data/' directory and all of its contents. This means that after running this command, files in the 'data/' directory can be modified, deleted, or renamed.

upvoted 3 times

 **linux_admin** 1 year, 10 months ago

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

upvoted 1 times

 **MissAllen** 2 years, 1 month ago

`cat > test.txt` is not reading a file, it is creating it. It allows you to type content in the shell and close the file with Ctrl-D. So, immutable attribute removal is appropriate. C is correct.

upvoted 3 times

 **TheRealManish** 2 years, 2 months ago

C seems wrong. This seems like a permissions question. we can not even view a file, so removing the immutable bit does not seem relevant.

upvoted 1 times

 **Veteran903** 2 years, 1 month ago

Not at all, he tried everything on the file and got denied, C is the right answer
upvoted 3 times

  **TheRealManish** 2 years, 1 month ago

thanks, i agree now

upvoted 1 times

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

Suggested Answer: B

linux_admin **Highly Voted** 1 year, 10 months ago

The correct command to ensure that both Java 7 and Java 8 are locally available for developers to use when deploying containers is:

B. docker image pull java:7

The "docker image pull" command is used to pull images from a Docker registry, such as Docker Hub, to the local system. The "java:7" option specifies the image to be pulled, in this case the Java 7 image.

So, in this example, the command "docker image pull java:7" pulls the Java 7 image to the local system, making it available for use when deploying containers.

upvoted 5 times

ericsrz **Most Recent** 1 year ago

The correct command to ensure that Java 7 is locally available for developers to use when deploying containers is:

B. docker image pull java:7

Here's why:

The docker image pull command is used to pull an image or a repository from a registry. This command will download the Java 7 image from the Docker Hub, making it locally available on the system.

upvoted 3 times

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

Suggested Answer: A

  **linux_admin**  1 year, 10 months ago

A. Cloud-init

Cloud-init is a popular tool used for provisioning virtual machines (VMs) in the cloud. It is commonly used to automate the initial configuration of cloud instances, such as setting up SSH keys, creating users, installing packages, and so on.

When installing packages during VM provisioning, cloud-init can be configured to run a script that installs the required packages, or to use a package management tool such as apt-get or yum to install the packages.

So, in this scenario, the cloud engineer should use cloud-init to accomplish the task of installing packages during VM provisioning. This will ensure that the packages are installed consistently and correctly, and that the VMs are properly configured for use.

upvoted 7 times

  **ericsrz**  1 year ago

The correct tool for a cloud engineer to use when installing packages during VM provisioning is:

A. Cloud-init

Here's why:

Cloud-init is a widely used approach to initialize cloud instances. It is a package that contains utilities for early boot configuration of a cloud instance. It can handle tasks such as setting up users, SSH keys, and installing packages, making it an ideal choice for this task.

upvoted 4 times

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.aws.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

Suggested Answer: D

Community vote distribution

C (100%)

 **linux_admin** Highly Voted 2 years, 4 months ago

Selected Answer: C

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.aws.ec2_instance" module, to create a cloud-based server with a public IP address. The "community.aws.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.
upvoted 7 times

 **Aj26a** Most Recent 1 year ago

Selected Answer: C

The code snippet provided is written in YAML and is used to define the creation of a cloud-based server instance. The syntax and structure are indicative of a playbook used by a configuration management and automation tool.

The correct technology that the systems administrator used to complete this task is:

C. Ansible

This can be determined by the format of the playbook and the use of modules such as community.aws.ec2_instance which are commonly used in Ansible for managing AWS resources.

So, the correct answer is C. Ansible.

upvoted 4 times

🗨️ 👤 **Alizadeh** 1 year, 10 months ago

The correct answer is C. Ansible

upvoted 2 times

🗨️ 👤 **Nvoid** 2 years, 7 months ago

Selected Answer: C

why can't the files just start with #!/dev/ansible !?!?

upvoted 2 times

🗨️ 👤 **Nvoid** 2 years, 7 months ago

or "#!/bin/ansible" ??

upvoted 1 times

🗨️ 👤 **MrGyKz** 2 years, 7 months ago

Selected Answer: C

as i use ansible , it looks like an ansible playbook

upvoted 2 times

🗨️ 👤 **bjornborg** 2 years, 7 months ago

Selected Answer: C

Terraform works with JSON, not YAML ... Ansible is YAML.

JSON uses { } YAML uses - and :

upvoted 2 times

🗨️ 👤 **drewbaby** 1 year, 7 months ago

While Ansible uses YAML, Terraform uses HCL (Hashicorp Configuration Language), not JSON.

upvoted 1 times

🗨️ 👤 **TheRealManish** 2 years, 8 months ago

Selected Answer: C

This looks more like ansible than terraform? I was googling ansible and the syntax looks like an ansible playbook and not a terraform config file.

upvoted 3 times

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`

Suggested Answer: B

Community vote distribution

B (100%)

🗨️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: B

The correct answer is B. `restorecon -R -v /var/www/html`.
upvoted 3 times

🗨️ 👤 **post20** 1 year, 3 months ago

Option 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.
upvoted 4 times

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.

Suggested Answer: B

Community vote distribution

B (100%)

🗨️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: B

The correct answer is B. Set PermitRootLogin to no in the /etc/ssh/sshd_config file.
upvoted 2 times

🗨️ 👤 **linux_admin** 1 year, 4 months ago

The administrator should set the PermitRootLogin option to no in the /etc/ssh/sshd_config file. This option determines whether the root user is allowed to log in to the SSH service remotely. Setting this option to no disables remote administrative login via SSH for the root user, providing additional security to the server.
upvoted 3 times

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.

Suggested Answer: A

Community vote distribution

A (67%)

C (33%)

🗳️ 👤 **Nvoid** Highly Voted 2 years, 7 months ago

I hope i never see this question..

upvoted 13 times

🗳️ 👤 **linux_admin** Highly Voted 2 years, 4 months ago

Selected Answer: A

Option A is correct. A bootloader is a software that initializes the system and starts the operating system. The bootloader is responsible for initializing the system, setting up the environment, and loading the operating system into memory.

Option B is incorrect. The bootloader does not mount the root filesystem. Mounting the root filesystem is usually done by the operating system after it has been loaded into memory by the bootloader.

Option C is correct to some extent. The bootloader does help to load the different kernels to initiate the OS startup process, but it's not limited to only kernels. The bootloader also has the ability to load other operating systems or other software.

Option D is incorrect. The bootloader does not trigger the start of all system services. Starting system services is usually done by the operating system after it has been loaded into memory by the bootloader.

upvoted 5 times

🗳️ 👤 **linux_admin** 2 years, 4 months ago

Discard The correct answer is C.

upvoted 1 times

🗳️ 👤 **JSHack** Most Recent 6 months, 1 week ago

Selected Answer: C

Answer: C. The bootloader's primary function is to locate and load the kernel.

Wrong:

- A. Device initialization is handled by the kernel during the early stages of the OS startup, not by the bootloader.
- B. The kernel is responsible for mounting the root filesystem once it has been loaded by the bootloader.
- D. Starting system services is handled by the init system after the kernel has booted and the root filesystem is mounted.

upvoted 2 times

🗳️ 👤 **RayRay2** 6 months, 3 weeks ago

Selected Answer: C

The primary function of a bootloader is:

- C. It helps to load the different kernels to initiate the OS startup process.

A bootloader is a small program that runs before the operating system starts. It is responsible for loading the operating system's kernel into memory and starting it.

Here's a brief rundown of the other options:

A: While some initialization of devices can occur, it's not the primary role of the bootloader.

B: Mounting the root filesystem typically happens later in the boot process, after the kernel is loaded.

D: The start of system services is managed by the init system (like systemd), not the bootloader.

upvoted 2 times

🗳️ 👤 **insanegrizly** 7 months, 4 weeks ago

Selected Answer: C

The bootloader is a low-level program that loads the operating system's kernel into memory and starts the boot process. It provides the user with options for selecting different kernels or operating systems (in a dual-boot environment, for example). It also passes control to the selected kernel to initiate the OS startup.

upvoted 2 times

🗳️ 👤 **Aj26a** 1 year ago

Selected Answer: C

The primary function of a bootloader is to load the operating system kernel and initiate the OS startup process. Specifically, it is responsible for:

C. It helps to load the different kernels to initiate the OS startup process.

The bootloader, such as GRUB (GRand Unified Bootloader), is the first software that runs when a computer starts. It typically provides a menu to select from multiple operating systems or kernels, loads the selected kernel into memory, and transfers control to the kernel.

So, the correct answer is C. It helps to load the different kernels to initiate the OS startup process.

upvoted 4 times

🗳️ 👤 **e418137** 1 year, 4 months ago

Selected Answer: A

A is correct. C is incorrect. There are not "different kernels to initiate the OS startup process." There's just one at a time. Obviously, the bootloader discovers and initializes the hardware that that kernel needs to use, or the OS would be useless. (The BIOS/EFI does a POST that includes hardware detection, but that's not enough for the kernel. The kernel needs to identify, configure, select drivers, etc.)

upvoted 1 times

🗳️ 👤 **Damon54** 1 year, 9 months ago

A and C are both correct 50%

upvoted 1 times

🗳️ 👤 **Damon54** 1 year, 9 months ago

Selected Answer: A

A. It initializes all the devices that are required to load the OS.

The bootloader plays a crucial role in the boot process of an operating system. One of its primary functions is to initialize the hardware devices that are necessary to load the operating system. This includes tasks such as identifying and configuring the CPU, memory, storage devices, and other essential hardware components to ensure that the system is in a suitable state to load the OS.

upvoted 2 times

🗳️ 👤 **KnifeClown1** 2 years, 4 months ago

Selected Answer: C

The correct answer is "C. It helps to load the different kernels to initiate the OS startup process."

A bootloader is a program that runs before the operating system (OS) starts. Its main function is to load the OS into memory and initiate the OS startup process. This involves loading the kernel, which is the core part of the OS that manages hardware resources and provides a platform for running user applications. The bootloader helps to load different kernels, which can be useful for testing and recovery purposes.

The other options listed are not functions of the bootloader:

Option A: Initializing devices is performed by the BIOS or UEFI (depending on the system), not the bootloader.

Option B: Mounting the root filesystem is performed by the kernel, not the bootloader.

Option D: Triggering the start of system services is performed by the OS, not the bootloader.

upvoted 3 times

🗳️ 👤 **MissAllen** 2 years, 7 months ago

I would go with C. It is the responsibility of the boot loader to load the kernel, but the responsibility of the kernel to load the root file system and initialize devices.

upvoted 2 times

  **TheRealManish** 2 years, 8 months ago

Terrible question

Could be:

A: yes it does initialize the hard drive to load the kernel

B: yes it mounts the filesystem to get the boot files to load the OS

C:Yes, it can load different kernels to initiate the OS startup process

upvoted 3 times

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

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

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.

Suggested Answer: D

Community vote distribution

D (100%)

  **[Removed]** 10 months, 3 weeks ago

Selected Answer: D

firewalld needs to be enabled so that it starts up on reboot

upvoted 2 times

  **TheRealManish** 1 year, 1 month ago

Selected Answer: D

Firewalld is disabled. Tested on my machine. when i create a rule in firewalld it does not show up on iptables output. I have both services running so no conflict.

upvoted 2 times

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>`

Suggested Answer: D

Community vote distribution

D (100%)

linux_admin **Highly Voted** 1 year, 4 months ago

Selected Answer: D

D. `docker inspect <container_id>`

The `docker inspect` command provides detailed information about a Docker container, including its configuration and state. Among the information available from this command is the `entrypoint`, which is the first command that is executed inside the container when it starts. By inspecting the container with `docker inspect`, the administrator can determine the first command that is executed inside the container and potentially diagnose why it is unable to start properly.

upvoted 5 times

Alizadeh **Most Recent** 10 months, 2 weeks ago

Selected Answer: D

The correct answer is D. `docker inspect <container_id>`.

upvoted 1 times

Nvoid 1 year, 7 months ago

Selected Answer: D

D is correct here, you use the `inspect` and default template which will output the `docker run` command.

upvoted 3 times

TheRealManish 1 year, 7 months ago

Is it really `docker inspect`? when i run `docker inspect` i just get a bunch of info about the container, but nothing about commands executed.

upvoted 1 times

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.

Suggested Answer: C

  **RayRay2** 6 months ago

Selected Answer: A

The correct answer is A. The checkdiskspace.timerunit should be enabled via systemctl.

Here's why:

In order for a systemd timer unit to be active and start the associated service, it needs to be enabled first. Enabling the timer makes sure it will be started at boot or when the systemd daemon reloads its configuration.

The other options are less likely to be the issue:

- B. The timers.targetshould be reloaded to get the new configuration: While reloading might be necessary after making changes, it won't start a timer that hasn't been enabled.
- C. The checkdiskspace.timersshould be configured to allow manual starts: This is not necessary for timers; timers run based on their scheduling, not manual starts.
- D. The checkdiskspace.timersshould be started using the sudo command: Starting a timer usually requires administrative privileges, but the primary issue would be that the timer needs to be enabled.



upvoted 2 times

  **Lorello2023** 6 months, 3 weeks ago

Selected Answer: A


the question clearly say the intent is not to start the script manually

upvoted 2 times

  **MissAllen** 7 months, 3 weeks ago

C is correct. Even root cannot start the service if the config file prevents it.

upvoted 2 times

  **NastyNutsu** 4 months, 1 week ago

C contradicts with the requirement, "The Linux Administrator does not want users to able to start the job"

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

`sudo systemctl enable checkdiskspace.timer`

`sudo systemctl start checkdiskspace.timer`

enabling the timer will ensure it runs according to the schedule without requiring manual intervention.

upvoted 1 times

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

Suggested Answer: D

🗨️ 👤 **Eikan** 6 months, 2 weeks ago

Selected Answer: D

V for vendend-- no, -V wget for verify.

upvoted 1 times

🗨️ 👤 **linux_admin** 10 months, 3 weeks ago

D. rpm -V wget

The rpm -V (or "verify") command can be used to check the integrity of installed RPM packages. The wget package, as well as all its files, can be verified using this command. If any files have been altered since the package was installed, rpm -V will report the differences. This can help the administrator determine whether any unauthorized changes have been made to the package, or if the files have been damaged in some way.

upvoted 4 times

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

Suggested Answer: B

  **angellorv** 9 months ago

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

upvoted 3 times

  **linux_admin** 10 months, 3 weeks ago

B. dig @10.10.10.20 example.com mx

The dig command is a useful tool for querying DNS servers to retrieve information about domain names.

upvoted 4 times

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.

```
# lsattr /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`

Suggested Answer: A

Community vote distribution

A (100%)

 **linux_admin** Highly Voted 10 months, 3 weeks ago

Selected Answer: A

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.

upvoted 7 times

 **Huckleberry** Most Recent 10 months, 3 weeks ago

should be 'lsattr' not 'lsattz'

upvoted 3 times

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`

Suggested Answer: A

Community vote distribution

A (100%)

🗉 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: A

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

upvoted 3 times

🗉 👤 **linux_admin** 1 year, 4 months ago

A. `systemctl status systemd-resolved.service`

The `systemctl status` command is used to show the current status of a system service in a Linux system that uses `systemd`. The status information includes the name of the service, its state (e.g. running, stopped), and any related error messages or warnings.

upvoted 2 times

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.

CentOS Command Prompt

```
[root@centos7] #
```

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

linux_admin Highly Voted 1 year, 4 months ago

`sudo yum install httpd``systemctl start httpd``systemctl enable httpd``systemctl status httpd`

upvoted 15 times

tutita 1 year, 1 month ago

correct, but no need to sudo since its already on root

upvoted 6 times

MissAllen Highly Voted 1 year, 7 months ago

Not sure why we need all those commands to install Apache and confirm it is running with the defaults. I would just do:

`yum install httpd``systemctl enable httpd`


```
systemctl start httpd
systemctl status httpd
upvoted 6 times
```

  **Lwarder1** 1 year, 4 months ago

Actually after looking again you have the right commands but wrong order:

```
yum install httpd
systemctl start httpd
systemctl enable httpd
systemctl status httpd
upvoted 5 times
```

  **Damon54** Most Recent 11 months, 1 week ago

```
dnf install -y httpd
```

Using a combination of cat and grep, confirm Apache's listen 80 value.

```
cat /etc/httpd/conf/httpd.conf | grep -i "listen 80"
```

Start and enable the httpd service by using systemctl and the appropriate subcommands provided below:

```
systemctl start httpd
systemctl enable httpd
systemctl is-active httpd
```

Enter the following command to configure the firewall to permit HTTP traffic:

```
firewall-cmd --zone=public --add-service=http --permanent
firewall-cmd --reload
```

upvoted 1 times

  **Lwarder1** 1 year, 5 months ago

Because you are correct. This practical was on xk0-004 and your explanation is correct.

upvoted 2 times

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

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: D

The correct answer is D. rpm -e.

upvoted 1 times

🗳️ 👤 **Lwarder1** 1 year, 4 months ago

-e flag is for erase

upvoted 3 times

🗳️ 👤 **linux_admin** 1 year, 4 months ago

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

upvoted 3 times

🗳️ 👤 **linux_admin** 1 year, 4 months ago

D. rpm -e

The rpm -e option is used to remove software packages from a Linux system that uses the RPM Package Manager. This option allows the administrator to remove one or more installed packages, along with all their dependencies.

upvoted 3 times



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

Suggested Answer: A

 **linux_admin**  10 months, 3 weeks ago

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.

upvoted 8 times

 **RayRay2**  6 months ago

Selected Answer: A

The correct answer is A. fsck.ext4 /dev/sda1.

Here's why:

The fsck.ext4 command is used to check and repair an ext4 filesystem. When a Linux system fails to start and provides an error message related to filesystem issues, running fsck.ext4 /dev/sda1 will check the integrity of the filesystem on /dev/sda1 and attempt to fix any detected errors.

Here's a brief explanation of the other options:

B. partprobe /dev/sda1: This command informs the operating system of partition table changes but does not check or repair filesystems.

C. fdisk /dev/sda1: This command is used for partitioning disks, not for checking or repairing filesystems.

D. mkfs.ext4 /dev/sda1: This command creates a new ext4 filesystem on the specified partition, which would erase all existing data on /dev/sda1, so it is not suitable for repair purposes.

upvoted 1 times

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

Suggested Answer: C

Community vote distribution



TheRealManish Highly Voted 2 years, 1 month ago

Selected Answer: A

This answer is wrong, cat out your /etc/login.defs file. those values are in there
upvoted 6 times

BryanSME Most Recent 1 year ago

linux_admin has it right: A is definitely correct, it's all right here, a quick look into the file shows below:
cat /etc/login.defs yields:

```
...
# Min/max values for automatic uid selection in useradd
#
UID_MIN 1000
UID_MAX 60000
# System accounts
SYS_UID_MIN 201
SYS_UID_MAX 999
```

```
#
# Min/max values for automatic gid selection in groupadd
#
GID_MIN 1000
GID_MAX 60000
# System accounts
SYS_GID_MIN 201
SYS_GID_MAX 999
```

upvoted 2 times

Tricee 1 year, 1 month ago

I can definitely understand why etc/default/useradd would be the answer. The question specifically stated that the users and groups were added "by default".

upvoted 1 times

🗨️ 👤 **linux_admin** 1 year, 10 months ago

Selected Answer: A

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:

upvoted 3 times

🗨️ 👤 **linux_admin** 1 year, 10 months ago

HOME /home/users/

The /etc/login.defs file is the correct configuration file to make changes to the default values for new users and groups. Other configuration files, such as /etc/security/limits.conf (used to set limits on resource utilization), /etc/default/useradd (used to set default options for the useradd command), or /etc/profile (used to set environment variables for users), are not used for this purpose.

upvoted 3 times

🗨️ 👤 **alimakkaya** 1 year, 11 months ago

Selected Answer: C

Question asks about newly created users . The file to be edited should be /etc/adduser.conf. And below is the section to be edited.

FIRST_[GU]ID to LAST_[GU]ID inclusive is the range of UIDs of dynamically

allocated user accounts/groups.

FIRST_UID=1000

LAST_UID=29999

FIRST_GID=1000

LAST_GID=29999

upvoted 1 times

🗨️ 👤 **alimakkaya** 1 year, 11 months ago

never mind . just delete the post, I got it wrong.

upvoted 1 times

🗨️ 👤 **Nvoid** 2 years, 1 month ago

Selected Answer: A

It's A.

upvoted 1 times

🗨️ 👤 **Nvoid** 2 years, 1 month ago

Mod plz update the answer its A.

upvoted 2 times

🗨️ 👤 **Nvoid** 2 years, 1 month ago

mods still need to update!

upvoted 2 times

🗨️ 👤 **MrGykz** 2 years, 1 month ago

Selected Answer: A

<https://man7.org/linux/man-pages/man5/login.defs.5.html>

reading default page for login.defs u can find these values written up : " SYS_UID_MAX (number), SYS_UID_MIN (number) Range of user IDs used for the creation of system users by useradd or newusers. "

upvoted 1 times

🗨️ 👤 **Veteran903** 2 years, 1 month ago

I dont understand how this can be wrong, correct answer is A, please fix it!

upvoted 2 times

🗨️ 👤 **MissAllen** 2 years, 1 month ago



Sorry, I agree with answer A, /etc/login.defs.

upvoted 3 times

🗨️ 👤 **MissAllen** 2 years, 1 month ago

Agreed, answer C.

upvoted 1 times

  **Nvoid** 2 years, 1 month ago

it is not C, it's A!

upvoted 1 times

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

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

Given the following analysis:

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

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

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

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

Which of the following is causing the error message?

- A. The administrator is not using a highly privileged account.
- B. The filesystem is mounted with the wrong options.
- C. SELinux file context is denying the ACL changes.
- D. File attributes are preventing file modification.

Suggested Answer: D

Community vote distribution


D (75%)

C (25%)

 **POGActual** Highly Voted 2 years, 3 months ago

I think it is D. With the append only attribute activated, you can only add information to the file; you cannot change it. This includes deleting or changing the ACL.

upvoted 8 times

 **Qubert2** Most Recent 8 months, 2 weeks ago

The append-only attribute (set with `chattr +a`) restricts modifications to the file, including changes to its metadata, which includes ACLs.

upvoted 1 times

 **e418137** 1 year, 4 months ago

Selected Answer: D

D. "File attributes are preventing file modification." It is not C: "SELinux file context is denying the ACL changes." SELinux does not discriminate: if an attribute can be set, then it can be unset. It's true that SELinux could be configured to totally disallow or selectively allow access to 'chattr', but that is not its default configuration and the question has no hint in regard to a change in SELinux policy. This can be easily demonstrated on any system using SELinux.

`touch file`

`sudo chattr +a file`


`lsattr file`

`rm file`

`sudo chattr -a file`

`lsattr file`

upvoted 4 times

 **ericsrz** 1 year, 6 months ago

Selected Answer: C

The correct option is C. SELinux file context is denying the ACL changes. The error message "setfacl: data.txt: operation not permitted" is caused by the SELinux file context denying the ACL changes.

upvoted 1 times

🗨️ 👤 **tutita** 2 years, 1 month ago

Selected Answer: D

it has the -----a-- (append attr) hence you cant remove nor modify the file

upvoted 2 times

🗨️ 👤 **post20** 2 years, 3 months ago

C. SELinux file context is denying the ACL changes.

The file context of the file is set to user_home_t and since the file is located under the /home directory, it is most likely that the file is located in a user's home directory which is restricted by SELinux. SELinux provides Mandatory Access Control (MAC) to restrict the access of the files and processes in a Linux system.

upvoted 1 times

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`

Suggested Answer: C

 **linux_admin** Highly Voted 10 months, 3 weeks ago

C. `ls | cpio -ov > cloud.cpio`

The cpio command is used to create and extract archive files in the cpio format. The -o option is used to create an archive file, and the -v option provides verbose output during the creation process. This will pipe the output of the ls command (which lists the files in the current directory) into the cpio command, which will create a new archive file cloud.cpio containing all the listed files. The > symbol is used to redirect the output of the cpio command to the cloud.cpio file, overwriting any existing file with the same name.

upvoted 5 times

 **ckl22** Most Recent 1 year ago

C is correct, STDOUT from ls (current directory) into the -ov (create archive and display files processed by cpio) and redirecting into the filename name cloud.cpio

upvoted 3 times

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`

Suggested Answer: A

 **linux_admin** Highly Voted 10 months, 3 weeks ago

A. `source ~/.bashrc`

The `source` command is used to re-read a shell configuration file and make its changes effective immediately in the current shell session. In this case, the administrator made changes to the `~/.bashrc` file and added an alias command, but the alias is not working. To make the changes in the `~/.bashrc` file effective, the administrator should run the following command:

```
source ~/.bashrc
```

This will cause the shell to re-read the `~/.bashrc` file and apply the changes, including the new alias. After running this command, the administrator should be able to use the alias as expected.

Option B is incorrect because the `read` command is used to read input from the user, not to re-read a configuration file. Option C is incorrect because the `touch` command is used to update the modification time of a file, not to re-read a configuration file. Option D is incorrect because the `echo` command is used to display text on the screen, not to re-read a configuration file.

upvoted 6 times

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

Suggested Answer: D

🗨️ 👤 **linux_admin** 10 months, 3 weeks ago

D. id_rsa.pub

The public authentication key, id_rsa.pub, is typically used to set up passwordless login, also known as SSH key-based authentication. In this scenario, the junior administrator has generated public and private authentication keys, and they need to be moved to the remote servers to set up passwordless login.

The public key, id_rsa.pub, is usually copied to the remote server and added to the ~/.ssh/authorized_keys file on the remote server. This allows the local system to authenticate with the remote server using the private key, id_rsa, without requiring a password.

upvoted 4 times

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

Device mismatch detected

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


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

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

- A. mount disk by device-id
- B. fsck -A
- C. mount disk by-label
- D. mount disk by-blkid

Suggested Answer: D

Community vote distribution



Nvoid Highly Voted 2 years, 7 months ago

Selected Answer: C

Changed yet again, but i'm serious this time. IT's C!!
upvoted 9 times

Misterymigi Most Recent 2 months, 2 weeks ago

Selected Answer: D

I think I'm going with D here.
Label can always change whereas blkid (yes a command) but retrieves the UUID and the UUID are unique to each filesystem.
upvoted 1 times

Misterymigi 2 months, 1 week ago

plus blkid shows label as well, its an all-in-one
upvoted 1 times

Storcaks 3 weeks, 1 day ago

It's a trick question. Mounting by by-blkid doesn't exist, only by-label is valid.

On almalinux 8

ls /dev/disk/

by-id by-label by-partuuid by-path by-uuid

It has to be C since only by-label exist of the options.

upvoted 1 times

insanegrizly 7 months, 4 weeks ago

Selected Answer: D

It's clear it's either C or D....

The best option is to mount to UUID which blkid would display, that's how you figure out the UUID, therefor I think that's why you you mount it to blkid....but it's entirely how you interpret the question for which answer you come up with.

upvoted 1 times

🗨️ **zionjr** 1 year, 1 month ago

D blkid: This command displays block device information, including a unique identifier called the block device ID. This ID remains constant for a specific disk regardless of device name assignment.

upvoted 1 times

🗨️ **IFBBPROSALCEDO** 1 year, 1 month ago

Selected Answer: D

To resolve the device mismatch issue and mount the disk correctly, the administrator should mount the disk using its UUID (Universally Unique Identifier). UUIDs are unique to each filesystem and do not change even if the device name (e.g., /dev/sdb, /dev/sdc) changes, making them ideal for consistent identification of disks. C. mount disk by-label: Disk labels can be useful but are not as unique and reliable as UUIDs. If labels are not unique, they can cause conflicts.

upvoted 1 times

🗨️ **e418137** 1 year, 4 months ago

Selected Answer: C

C. By lable. Why? (A) There is no "device-id." (B) The 'fsck' command repairs file systems. (D) The 'blkid' command shows block device attributes. (The sane choice is UUID because the UUID is embedded in the file system of the block device, but that's not an option to choose.)

upvoted 2 times

🗨️ **DRVision** 1 year, 6 months ago

Selected Answer: D

Option D, mount disk by-blkid, is generally better than option C, mount disk by-label, for a few reasons:

Uniqueness: The block ID (blkid) is a unique identifier for each block device, which is guaranteed to be the same across reboots and even if the disk is moved to a different machine. On the other hand, disk labels are not guaranteed to be unique and can be changed by the user, leading to potential conflicts.

Presence: Not all filesystems support labels, and even when they do, labels are not always set by default. This means that you might not be able to mount a disk by label if the label has not been set. In contrast, all block devices have a blkid.

Consistency: The blkid remains consistent even if the device name (e.g., /dev/sda, /dev/sdb) changes due to hardware changes or reconfiguration. Disk labels do not have this issue, but combined with the points above, using blkid is generally more reliable.

Therefore, while both methods can be used to mount a disk, using the blkid (option D) is typically more reliable and less prone to errors or conflicts.

upvoted 2 times

🗨️ **e418137** 1 year, 4 months ago

There is no block id.

upvoted 1 times

🗨️ **DRVision** 1 year, 6 months ago

When a server is cloned, the filesystem labels (mount disk by-label) are also cloned. This means that if both the original and cloned servers are on the same network, there could be two disks with the same label, leading to confusion and potential errors.

On the other hand, the block ID (blkid) is unique for each block device, even across clones. This means that even if a disk is cloned, the clone will have a different blkid. Therefore, mounting by blkid avoids the potential conflicts that can arise when cloning servers.

So, even in the case of server cloning, using the blkid (option D) is typically more reliable and less prone to errors or conflicts. It ensures that the correct disk is mounted, regardless of any changes in the hardware configuration or cloning of disks.

upvoted 1 times

🗨️ **e418137** 1 year, 4 months ago

There is no block id.

upvoted 1 times

🗨️ **wait4thebus** 1 year, 7 months ago

Selected Answer: C

See comment below. Can someone present a clearer argument on why block ID would be correct and not mount disk by-label?

upvoted 1 times

🗨️ **e418137** 1 year, 4 months ago

There is no block id.

upvoted 1 times

🗳️ 👤 **wait4thebus** 1 year, 7 months ago

I am going with mount disk by-label because the question states that the IDs were from a clone and the UUIDs are tied to the source device of the cloning. Since the UUIDs of the source of the clone are unique to just that source device and mean nothing to anything else, it would make sense to use the more generic mount disk by-label method. So, C would be the answer.

upvoted 2 times

🗳️ 👤 **linux_admin** 2 years, 4 months ago

Selected Answer: D

The "mount disk by-blkid" option allows the administrator to mount the disk using a unique identifier (the block ID) instead of the device name. This is particularly useful when cloning a server and the device names have changed, as it helps to ensure that the correct disk is being mounted.

The "mount disk by device-id" option is not correct because it does not exist in the standard Linux mount command.

The "fsck -A" option is used to check and repair file systems, but it does not mount the disk.

The "mount disk by-label" option is also incorrect as it mounts the disk based on the disk label, but if the disk label has changed as a result of the clone, this option will not work.

upvoted 1 times

🗳️ 👤 **linux_admin** 2 years, 4 months ago

Discard

upvoted 2 times

🗳️ 👤 **e418137** 1 year, 4 months ago

There is no block id.

upvoted 1 times

🗳️ 👤 **lo_01234_ol** 2 years, 7 months ago

Selected Answer: C

Here ya' go: <https://unix.stackexchange.com/questions/644708/how-can-i-prevent-disk-uuid-mismatch-when-cloning-a-machine>

upvoted 3 times

🗳️ 👤 **Nvoid** 2 years, 7 months ago

Device label better, i'm actually going to have to do this once i start my job doing this. lol

upvoted 2 times

🗳️ 👤 **Nvoid** 2 years, 7 months ago

Selected Answer: A

I changed by vote to A - Device ID is better, Ref: https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html/managing_file_systems/assembly_overview-of-persistent-naming-attributes_managing-file-systems

upvoted 1 times

🗳️ 👤 **TheRealManish** 2 years, 7 months ago

I don't know man, your link talks about UUID. selection A is referring to "device-id" which is not in your link, and doesnt really show up on google as anything meaningful in regards to linux.

upvoted 1 times

🗳️ 👤 **TheRealManish** 2 years, 7 months ago

Selected Answer: C

According to this link the answer should be C: <https://unix.stackexchange.com/questions/644708/how-can-i-prevent-disk-uuid-mismatch-when-cloning-a-machine>

upvoted 4 times

🗳️ 👤 **TheRealManish** 2 years, 7 months ago

not feeling super confident here, does anyone else have thoughts?

upvoted 1 times

🗳️ 👤 **Nvoid** 2 years, 7 months ago

blkid is a command, i'm choosing C, i remember from the class i took Label is more reliable and easier to work with.

upvoted 2 times

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

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

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

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

Suggested Answer: B

Community vote distribution

B (100%)

 **linux_admin** Highly Voted 10 months, 3 weeks ago

B. `lscpu`

The `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.

upvoted 5 times

 **alimakkaya** Most Recent 11 months, 1 week ago

Selected Answer: B

`lscpu` will output following line for virtualization support.

...

Virtualization: VT-x

...


upvoted 2 times

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+rxw /project/access2all`

Suggested Answer: A

Community vote distribution



IFBBPROSALCEDO 7 months ago

Selected Answer: A

This sets the sticky bit on the directory `/project/access2all`. The sticky bit ensures that only the owner of a file can delete or rename it within that directory, preventing non-owners from deleting or modifying files they do not own.

upvoted 2 times

e418137 10 months, 3 weeks ago

Selected Answer: A

A. Just look for yourself: `'ls -ld /tmp'`. This directory accomplishes the goal of the question. That was done in octal `'chmod 1777 /tmp'` or in UGO shorthand `'chmod +t /tmp'`. 2770 is the setGID bit, explained by DRVision.

upvoted 2 times

DRVision 1 year ago

Selected Answer: A

Option A, `chmod +t /project/access2all`, is better than option C, `chmod 2770 /project/access2all`, for preventing the deletion or modification of files from non-owners because of the following reasons:

Sticky Bit: Option A sets the sticky bit on the directory. The sticky bit restricts deletion or renaming of files within the directory to the file owners, the directory owner, or the root user. This means even if a user has write permission to the directory, they cannot delete or rename files owned by others.

Setgid and Permissions: Option C sets the setgid bit and gives the group read, write, and execute permissions. The setgid bit causes new files and directories created in the directory to inherit the group ownership of the directory, but it does not prevent file deletion or modification by non-owners. The permissions 770 give the owner and the group full permissions (read, write, execute), but they do not prevent file deletion or modification by non-owners who are in the group.

Therefore, to specifically prevent deletion or modification of files from non-owners, option A is the better choice.

upvoted 3 times

Alizadeh 1 year, 4 months ago

Selected Answer: C

The correct answer is C. `chmod 2770 /project/access2all`.

upvoted 1 times

linux_admin 1 year, 10 months ago

Selected Answer: C



C. `chmod 2770 /project/access2all`.

This option sets the setgid (`chmod +2`) and the permissions `rxw` (`chmod 700`) on the directory. This means that files created within the directory will inherit the group ownership of the directory, rather than the user's primary group, and members of the group will have full permissions to read, write, and execute files within the directory. Non-group members will not have access to the directory. This helps to ensure that files created within the directory can only be modified by members of the group and not by non-owners.

upvoted 2 times

tutita 1 year, 7 months ago

correct, the sticky bit only works for not deleting a file, the question states modify and delete the files hence is option C. chmod 2770
upvoted 1 times

  **tutita** 1 year, 7 months ago

I meant to say, the sticky bit denies to other users the right to delete or RENAME a file. the question states "MODIFY and delete" for non-ownerS. Im guessing here by non ownerS means users that belong to a group. terrible wording but since its stating modify not rename Im going for C.

upvoted 1 times

  **linux_admin** 1 year, 10 months ago

Answer A (chmod +t /project/access2all) sets the sticky bit on the directory, but this only affects the ability of regular users to delete files. The sticky bit does not prevent regular users from modifying files within the directory. To prevent modification of files within the directory by non-owners, you would need to set more restrictive permissions using chmod or using access control lists (ACLs).

So, answer A is not the correct solution to the problem of preventing the modification or deletion of files within the directory by non-owners.

upvoted 1 times

  **BreakOff874** 1 year, 10 months ago

You are wrong. The sticky bit does prevent users, other than the owner, to delete files. The question is asking for a solution that will prevent non- owner(groups and others) from deleting files. chmod +t will do the job.

Answer C is giving rwx to groups (non-owner)

upvoted 5 times

  **Huckleberry** 1 year, 10 months ago

restricted deletion flag or sticky bit (t).

upvoted 2 times

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=1" >> /etc/sysctl.conf`

Suggested Answer: BC

Community vote distribution



🗳️ 👤 **Storcaks** 3 weeks, 1 day ago

Selected Answer: BC

Peopling saying CE:

You are assuming that C will be run before E. But if E is run before C then it will not be persistent. B & C will always work in any order. BC is a safer choice.

upvoted 1 times

🗳️ 👤 **Mistermiyagi** 2 months, 2 weeks ago

Selected Answer: BC

I kind of think BC and CE are both correct answers. They both do literally the same thing.

upvoted 1 times

🗳️ 👤 **HappyDay030303** 2 months, 3 weeks ago

Selected Answer: BC

- B. Enable it immediately (at runtime)
 - C. ensures the setting is applied every time the system boots
- upvoted 1 times

🗳️ 👤 **Qubert2** 8 months, 1 week ago

Selected Answer: CE

Correct answer is C and E. C changes the config file (which is read each time on boot) and E loads it immediately.

B (`sysctl -w net.ipv4.ip_forward=1`) make the change immediately but does not persist. it's wrong.

upvoted 1 times

🗳️ 👤 **Bimbo_12** 1 year, 4 months ago

To persistently enable IPv4 forwarding in a Linux system, you typically need to modify the system configuration files. The two commands you can actually use together to accomplish this task are:

C. `echo "net.ipv4.ip_forward=1" >> /etc/sysctl.conf`: This appends the configuration `net.ipv4.ip_forward=1` to the `/etc/sysctl.conf` file, which will enable IPv4 forwarding persistently across reboots.

E. `sysctl -p`: This command reloads the `sysctl` settings from the configuration files, including `/etc/sysctl.conf`, so that the changes take effect immediately without needing a system reboot.

So the correct combination of commands is C and E. They ensure that IPv4 forwarding is enabled persistently and immediately without requiring a reboot.

upvoted 2 times

🗳️ 👤 **e418137** 1 year, 4 months ago

Selected Answer: BC

B & C accomplish the same goal: write 'net.ipv4.ip_forward=1' to '/etc/sysctl.conf'. (And whatever is in '/etc/sysctl.conf' gets used persistently.) (A) prints the current value of 'net.ipv4.ip_forward'. (D) turns on IPv4 forwarding temporarily (e.g.: gone after reboot). (E) reads 'sysctl' configuration file(s). (F) relates to IPv6, not IPv4 in the question.

upvoted 2 times

🗳️ 👤 **bongobo** 1 year, 4 months ago

D. echo 1 > /proc/sys/net/ipv4/ip_forward

IS CORRECT ALSO

upvoted 1 times

🗳️ 👤 **p24** 6 months ago

but not persistent

upvoted 1 times

🗳️ 👤 **DRVision** 1 year, 6 months ago

Selected Answer: CE

C. echo "net.ipv4.ip_forward=1" >> /etc/sysctl.conf E. sysctl -p

The command echo "net.ipv4.ip_forward=1" >> /etc/sysctl.conf[C] is used to add the line net.ipv4.ip_forward=1 to the /etc/sysctl.conf file. This line enables IPv4 forwarding when the system boots.

The command sysctl -p[E] is used to reload the sysctl settings from the /etc/sysctl.conf file. This makes the changes take effect immediately without requiring a system reboot.

upvoted 2 times

🗳️ 👤 **DRVision** 1 year, 6 months ago

Option B, sysctl -w net.ipv4.ip_forward=1, does enable IPv4 forwarding, but it does not persist after a system reboot. This command changes the runtime settings of the system, which are reset to their default values when the system is rebooted.

On the other hand, Option C, echo "net.ipv4.ip_forward=1" >> /etc/sysctl.conf, writes the setting into the /etc/sysctl.conf file, which is read at boot time. This makes the change persistent across reboots.

Option E, sysctl -p, is used to reload the sysctl settings from the /etc/sysctl.conf file immediately, without needing to reboot the system. This is why options C and E are chosen for persistently enabling IPv4 forwarding.

upvoted 2 times

🗳️ 👤 **giomax** 1 year, 7 months ago

Selected Answer: BD

just tested B and D is correct

upvoted 1 times

🗳️ 👤 **salthedhash** 1 year, 7 months ago

Selected Answer: CE

C. echo "net.ipv4.ip_forward=1" >> /etc/sysctl.conf

This command adds the net.ipv4.ip_forward=1 line to the /etc/sysctl.conf file, which will be read during system startup.

E. sysctl -p

This command is used to apply changes from the sysctl configuration file, including the changes made in /etc/sysctl.conf. It makes sure that the changes take effect without requiring a system restart.

upvoted 1 times

🗳️ 👤 **Rob74613** 2 years, 1 month ago

Selected Answer: CE

Alot of people are putting BC and I thought I agreed with it until I looked closer at the question, and I found 2 things wrong with BC

1. B doesnt actually make it persistent across system reboots only C does that
2. The question asks what commands can be used together, seems odd to run these commands together

C and E however

C. will turn on ipv4 forwarding via appending it to the /etc/sysctl.conf file

D. will apply and reload the changes without a reboot of the system

upvoted 1 times

🗳️ 👤 **linux_admin** 2 years, 4 months ago

Selected Answer: CE

The commands that can be used together to persistently enable IPv4 forwarding in a Linux system are:

C. `echo "net.ipv4.ip_forward=1" >> /etc/sysctl.conf`

E. `sysctl -p`

The first command `echo "net.ipv4.ip_forward=1" >> /etc/sysctl.conf` is used to add the `net.ipv4.ip_forward=1` line to the end of the `/etc/sysctl.conf` file. This file is used to configure kernel parameters at runtime. The change made in this file will persist across reboots.

The second command `sysctl -p` is used to apply the changes made to the `/etc/sysctl.conf` file. It reloads the configuration file and sets the new values for the specified parameters.

The other options `sysctl net.ipv4.ip_forward`, `sysctl -w net.ipv4.ip_forward=1`, `echo 1 > /proc/sys/net/ipv4/ip_forward`, and `echo "net.ipv6.conf.all.forwarding=1" >> /etc/sysctl.conf` are not suitable for persistently enabling IPv4 forwarding, as the changes made with these commands will not persist after a reboot.

upvoted 3 times

  **linux_admin** 2 years, 4 months ago

Discard I'm choosing BC.

upvoted 2 times

  **KnifeClown1** 2 years, 4 months ago



Selected Answer: CE

The correct options to persistently enable IPv4 forwarding in Linux are C and E.

Option C adds the line `"net.ipv4.ip_forward=1"` to the file `/etc/sysctl.conf` which is used to configure kernel parameters at boot time. This ensures that IPv4 forwarding is enabled each time the system starts.

Option E is used to reload the `sysctl` configuration from the `/etc/sysctl.conf` file, which in turn sets the IPv4 forwarding value to 1 as specified in the configuration file.

upvoted 1 times

  **KnifeClown1** 2 years, 4 months ago

Correct answer = BC

upvoted 2 times

  **Pinnubhai** 2 years, 5 months ago

Selected Answer: BC

1.to allow persistent IPv4 packet forwarding: `sysctl -w net.ipv4.ip_forward=1`

2. Once the system is satisfactorily tuned, make the new values permanent by modifying `/etc/sysctl.conf` or the `/etc/sysctl.d/` directory.

upvoted 2 times

  **Ckl22** 2 years, 6 months ago

Selected Answer: BC

BC are the only ones that enable IPv4 forwarding persistently

upvoted 2 times

  **TheRealManish** 2 years, 7 months ago

Selected Answer: BC

Proof that it is B and C : <https://www.systutorials.com/setting-up-gateway-using-iptables-and-route-on-linux/>

upvoted 2 times

  **Nvoid** 2 years, 7 months ago

B & C are correct here.

B & D do the same thing, but B is a command way and more versatile.

C makes it persistently enabled as the question asks.

upvoted 1 times

  **clmason1994** 2 years, 7 months ago

I agree with B but I disagree with C. It should be D given the information

<https://linuxconfig.org/how-to-turn-on-off-ip-forwarding-in-linux>

upvoted 1 times

  **Ckl22** 2 years, 6 months ago

D is temporary, C is persistent through reboots

upvoted 1 times

🗨️ 👤 **MrGy kz** 2 years, 7 months ago

D talks about ipv6 not ipv4 , the task asks about ipv4

upvoted 1 times

🗨️ 👤 **MrGy kz** 2 years, 7 months ago

Even in your used link it states:

Using either method above will not make the change persistent. To make sure the new setting survives a reboot, you need to edit the /etc/sysctl.conf file.

```
# sudo nano /etc/sysctl.conf
```


upvoted 1 times

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`

Suggested Answer: A

Community vote distribution



linux_admin Highly Voted 1 year, 10 months ago

Selected Answer: C

The correct answer is C, "find /var/log -type f -mtime +180 -exec rm {} ;". This command will accomplish the task of finding and removing all log files that were modified more than 180 days ago.

Here's a breakdown of the command:

The find command is used to search for files and directories in a specific location, in this case, "/var/log".

The -type f option is used to search only for regular files, excluding directories and other file types.

The -mtime +180 option is used to select only files that have a modification time more than 180 days ago.

The -exec rm {} \; option is used to execute the rm command on each file that matches the criteria specified in the previous options. The curly braces {} are replaced with the name of each file found and the backslash \; is used to end the command.

upvoted 7 times

linux_admin 1 year, 10 months ago

The other options, A and D, contain incorrect syntax or incorrect commands. Option B only has the "find" and "rm" commands and is missing the options to specify the type of files to search for and the time frame for modification. Option D uses the incorrect option "-remove" instead of "-exec rm {} ;".

upvoted 2 times

Nvoid Highly Voted 2 years, 1 month ago

Selected Answer: C

C is Correct.

upvoted 7 times

bc1235813 Most Recent 9 months, 3 weeks ago

Selected Answer: C

"A" is going after directories older than 180 days, "B" & "D" are gibberish - "C" would be right if the semicolon was there at the end of the command.

upvoted 1 times

e418137 10 months, 3 weeks ago

Selected Answer: A

A. The 'find' command with its '-exec' option requires an escaped semi-colon, ';' (B) There is no predicate, modified. (C) This would work just as well as answer, A, if the command ended with a semi-colon. (D) Nonsense.

upvoted 1 times

BryanSME 1 year ago

THE CORRECT OPTION IS C. I've tested this, but used 1600 days so found only one file that old, used option C and it deleted it only, results below:

```
# find /var/log -type f -mtime +1600
```

```
/var/log/anaconda/syslog <<<<<<<<<< found this file over 1600 days old-didn't want to remove anything important
```

```
# find /var/log -type f -mtime +1600 -exec rm {} \; <<<<<<<<<< Executed option C
```

```
# find /var/log -type f -mtime +1600 <<<<<<<<<< Ran command again and /var/log/anaconda/syslog is no longer present
```

upvoted 3 times

🗨️ 👤 **Damon54** 1 year, 5 months ago

in question C ; end the command is not present , the command not work ! A is correct ?
upvoted 1 times

🗨️ 👤 **KnifeClown1** 1 year, 10 months ago

Selected Answer: C

The correct command to accomplish this task is C:

This command uses the find utility to search for files in the /var/log directory and its subdirectories. The -type f option specifies to search only for files (not directories), and the -mtime +180 option specifies to search for files that were modified more than 180 days ago. The -exec rm {} \; option specifies to execute the rm command on each file found by the find command. The {} characters are a placeholder for each file that is found, and the \; at the end of the command is used to terminate the -exec option.

Note that the other options listed are incorrect and should not be used.

upvoted 2 times

🗨️ 👤 **Ckl22** 2 years ago

Selected Answer: C

The answer is C

upvoted 4 times

🗨️ 👤 **MissAllen** 2 years, 1 month ago

Disagree. Correct answer should be C. The question refers to removing log "files". The type option for find needs -f for files.

upvoted 7 times

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
```

Suggested Answer: A

Community vote distribution

C (64%)

A (36%)

  **e418137** 10 months, 3 weeks ago

Selected Answer: C

They don't mention the purpose of the NAT, so keep it simple for the "entry level" exam. With the most common form of NAT or IP Masquerading, post-routing alters packets as they leave the system.

upvoted 1 times

  **BryanSME** 1 year ago

Option C. did execute without errors:

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

So I'm with LRISB, C is the correct answer

upvoted 3 times

  **DRVision** 1 year ago

Selected Answer: C

Option B is the correct one: `echo 1 > /proc/sys/net/ipv4/ip_forward; iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE`

This command does two things:

`echo 1 > /proc/sys/net/ipv4/ip_forward` enables IP forwarding, which is necessary for the server to forward packets between interfaces.

`iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE` sets up NAT (Network Address Translation) so that packets leaving the eth0 interface will have their source IP address replaced with the IP address of the eth0 interface (i.e., the IP address of the server). This is necessary for the server to act as a router.

Option A is incorrect because it uses PREROUTING instead of POSTROUTING. The PREROUTING chain is used for DNAT (Destination NAT), not for SNAT (Source NAT) which is what we want in this case.

upvoted 1 times

🗨️ 👤 **LKRISB** 1 year, 7 months ago

The iptables command "`-t nat -A POSTROUTING -o eth0 -j MASQUERADE`" configures NAT on the server. It adds a rule to the nat table in iptables that performs source NAT (SNAT) on outgoing packets. The "`-o eth0`" option specifies the outgoing interface (eth0 in this case), and "`-j MASQUERADE`" instructs iptables to modify the source IP address of outgoing packets to match the IP address of the interface, effectively masquerading the internal IP addresses.

Options a and d are incorrect because they use the PREROUTING chain in iptables, which is used for modifying packets as they enter the system. In the context of setting up a router, we need to modify packets as they leave the system, so we should use the POSTROUTING chain.

Option b is incorrect because it uses the "`-D`" flag, which stands for "delete," to remove a rule from iptables. However, in this case, we need to add a rule to configure NAT, not delete an existing rule.

Therefore, the correct answer is

upvoted 2 times

🗨️ 👤 **LKRISB** 1 year, 7 months ago

Selected answers is : C

upvoted 1 times

🗨️ 👤 **BreakOff874** 1 year, 10 months ago

Selected Answer: A

postrouting does not forwards incoming traffic.

upvoted 1 times

🗨️ 👤 **linux_admin** 1 year, 10 months ago

Selected Answer: C

C. `echo 1 > /proc/sys/net/ipv4/ip_forward`

`iptables -t nat -A PREROUTING -o eth0 -j MASQUERADE`

Explanation:

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.

upvoted 2 times

🗨️ 👤 **linux_admin** 1 year, 10 months ago

Option A (`echo 1 > /proc/sys/net/ipv4/ip_forward` and `iptables -t nat -A PREROUTING -i eth0 -j MASQUERADE`) is not preferred because it only enables IP forwarding and NAT (network address translation) for incoming traffic on the eth0 interface. This may not be sufficient for routing all the traffic from a remote site.

In a real-world scenario, the router would need to route traffic for both incoming and outgoing traffic. Option C (`echo 1 > /proc/sys/net/ipv4/ip_forward` and `iptables -t nat -A PREROUTING -o eth0 -j MASQUERADE`) is more complete as it enables IP forwarding and NAT for both incoming and outgoing traffic on the eth0 interface.

upvoted 1 times

🗨️ 👤 **Ckl22** 2 years ago

Selected Answer: A

I think the answer is A, as with IP forwarding and PREROUTING, both modifying the packet as it arrives from outside the private LAN, and altering the destination address to a preconfigured mapping

upvoted 2 times

🗨️ 👤 **Nvoid** 2 years, 1 month ago

Selected Answer: A

I picked `A` because its a "remote system" and taking in packets to route which is -i for the "in" interface and the "prerouting". And it's implying the there needs to be traffic between both remote sites.

upvoted 1 times

🗨️ 👤 **TheRealManish** 2 years, 1 month ago

i think C, please read this and let me know if you still think A after reading.. thanks

upvoted 1 times

🗨️ 👤 **TheRealManish** 2 years, 1 month ago

<https://www.adamintech.com/configure-nat-masquerading-in-iptables/>

upvoted 1 times

🗨️ 👤 **Nvoid** 2 years, 1 month ago

Thanks i read it, i believe you need both a prerouting rule and a postrouting rule:

prerouting is incoming packages.

postrouting is outgoing packets to other networks.

so -i should be used for "in" which would be prerouting.

and -o should be used for "out" which would be posting routing.

i'm sticking to A,

3 weeks ago someone made a comment that they got 60 out of 63 on the test, so most of the questions i'm thinking are correct, just something thats one the back of my mind.

upvoted 2 times

🗨️ 👤 **TheRealManish** 2 years, 1 month ago

Selected Answer: C

The more i research this one, the more the answer is C. Several links agree.

upvoted 4 times

🗨️ 👤 **TheRealManish** 2 years, 1 month ago

I don't know the answer for sure, but this link makes it seem like it could be C? <https://bobcares.com/blog/iptables-nat-masquerade/>

upvoted 4 times

🗨️ 👤 **Nvoid** 2 years, 1 month ago

its A or C 50/50 chance.

upvoted 2 times

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

```
Oct 20 03:45:50 hostname kernel: iptables denied: IN=eth0 OUT=
MAC=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx: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 URG=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: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 URG=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: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 URG=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: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 URG=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
```

Suggested Answer: B

Community vote distribution

B (100%)

 **IFBBPROSALCEDO** 7 months ago

Selected Answer: B

To resolve the issue where incoming connections to ports 80 and 22 are being denied, you need to add rules to the INPUT chain of iptables to accept TCP traffic on these ports. This will allow the firewall to accept incoming HTTP and SSH connections, thus restoring access to the organization's web page and other services

upvoted 3 times

 **linux_admin** 1 year, 10 months ago

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.

upvoted 4 times

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`

Suggested Answer: B

Community vote distribution

B (100%)

🗨️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: B

The correct answer is B. `mv .ssh/authorized_key .ssh/authorized_keys`
upvoted 1 times

🗨️ 👤 **linux_admin** 1 year, 4 months ago

B. `mv .ssh/authorized_key .ssh/authorized_keys`

This 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.

upvoted 4 times

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.

Suggested Answer: A

Community vote distribution

A (83%)

C (17%)

  **e418137** 10 months, 3 weeks ago

Selected Answer: A

A. RB34561 explains it. (Tangentially, the file can be deleted by its owner and the root user. The file's permissions can be changed by its owner and the root user.)

upvoted 2 times


  **RB34561** 1 year ago

Selected Answer: A

Selecting A -

The output from the `ls -l` command shows the permissions of the startup file as a series of dashes (-----), this means that there are no permissions set for the owner, group, or others. This would prevent any user, including the root user, from reading, writing, or executing the file.

upvoted 4 times

  **Damon54** 1 year, 3 months ago

Selected Answer: C

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.

upvoted 1 times

  **JRS99** 6 months ago

"Shouldn't" and "cannot" are definitely different. We can all see the file does indeed belong to root, so C would be incorrect.

upvoted 1 times

  **JRS99** 6 months ago

"Shouldn't" and "cannot" are definitely different. We can all see the file does indeed belong to root, so C would be incorrect.

upvoted 1 times

  **angellorv** 1 year, 9 months ago

-l flag display file permissions


owner (root) and group (root) have no permissions

upvoted 1 times

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`

Suggested Answer: D

 **linux_admin** Highly Voted 1 year, 10 months ago
D. `chmod 1755 devops`

The sticky bit is a special permission that can be set on directories, and it is represented by the number "1" in the binary representation of the file permission mode. To set the sticky bit and keep the 755 file permission, you would set the permission mode to 1755. The "`chmod 1755 devops`" command sets the file permission mode to 1755, which includes the sticky bit and the 755 file permissions.
upvoted 7 times

 **e418137** Most Recent 10 months, 3 weeks ago
Also can be accomplished with ``chmod +t directory_name``.
upvoted 4 times

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.

Suggested Answer: C

 **linux_admin** Highly Voted 10 months, 2 weeks ago

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

When a Linux system is booted, it starts a target, which is a group of units that define the system state. The administrator ran the command "`systemctl isolate graphical.target`" which switches the system to the graphical target and provides a GUI interface. However, if the default target is not set to `graphical.target`, the system will revert back to the non-GUI terminal after a reboot.

upvoted 6 times

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.

Suggested Answer: A

Community vote distribution



Nvoid Highly Voted 2 years, 1 month ago

Selected Answer: C

I'm picking C, this isn't a MySQL database test so knowing what a semaphores seems outside of the scope of linux+.

upvoted 6 times

Damon54 Most Recent 1 year, 3 months ago

Answer C could be a consequence or symptom of the out-of-memory problem, but it is not the direct cause. Therefore, answer D is much more likely to be the main explanation for the problem.

upvoted 2 times

Damon54 1 year, 3 months ago

Selected Answer: D

Why not D ?

D. The amount of RAM allocated to the server is too high.

upvoted 1 times

comptiamac1 8 months, 1 week ago

Please read once again: "amount of RAM allocated to the server", not "used by the server". Key term: "ALLOCATED". It would never be a problem/issue. It's like saying, that you gave 20T drive to the server and this caused issue. More memory = better performance.

upvoted 4 times

linux_admin 1 year, 10 months ago

Selected Answer: C

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.

upvoted 2 times

linux_admin 1 year, 10 months ago

Option A "The process mysqld is using too many semaphores" is not relevant to the issue described in the log message. Semaphores are a type of synchronization mechanism that are used to control access to shared resources in a multithreaded environment. They are implemented as a type of lock that can be acquired and released by multiple processes.

While semaphores can be a source of performance issues and can cause problems in a system, they do not have a direct impact on memory usage. If a process is using too many semaphores, it may cause other issues such as deadlocks or performance degradation, but it will not cause the system to run out of memory.

In this case, the log message indicates that the process mysqld was terminated by the Linux kernel due to a lack of available memory, not due to an issue with semaphores. The message specifically mentions "out of memory", which suggests that the root cause of the issue is related to memory, not semaphores.

upvoted 2 times

🗨️ 👤 **KnifeClown1** 1 year, 10 months ago

Selected Answer: A

The message in /var/log/messages is indicating that the process mysqld is using too many semaphores, which is causing the connection issue. So the correct answer is A: The process mysqld is using too many semaphores.

Semaphores are used to control access to shared resources in a system. When a process requires access to a shared resource, it acquires a semaphore. When the process has finished using the resource, it releases the semaphore. If a process requires too many semaphores, it can cause other processes to block, leading to performance issues. In this case, the mysqld process is using too many semaphores, causing connections to the MariaDB service to be closed unexpectedly.

It's important to note that other factors such as memory, disk space, CPU usage, etc. should also be monitored to determine the root cause of the issue.

upvoted 1 times

🗨️ 👤 **KnifeClown1** 1 year, 10 months ago

Correct answer C:

upvoted 1 times

🗨️ 👤 **TheRealManish** 2 years, 1 month ago

Selected Answer: C

Maybe I'm wrong, but semaphore is just a non-negative shared variable? This just seems like the kernel was starved for memory and killed the process.. C, something is starving the server resources.

upvoted 4 times

🗨️ 👤 **Veteran903** 2 years, 1 month ago

I took a second look at this question, the log message clearly indicates with parenthesis the issue is within mysqld and just for that reason I'm gonna change my answer to A, CompTIA is very tricky with the way they word their questions, C can be an answer but not in this particular question.

upvoted 1 times

🗨️ 👤 **TheRealManish** 2 years, 1 month ago

Hey, thanks. I'm having a real hard time finding anything in google resembling this answer. can you possibly elaborate? thanks so much.. unbelievable that something so hard to find on google is an intermediate level cert studies.

upvoted 1 times

🗨️ 👤 **Veteran903** 2 years, 1 month ago

I'm with you, answer is C

upvoted 2 times

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

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **linux_admin** 1 year, 10 months ago

D. \$DISPLAY

The DISPLAY environment variable tells the X Window System which display to use for any graphical applications that are launched. When using X11 forwarding over SSH, the DISPLAY variable must be set in the remote shell to the value of the display on the local system. This allows graphical applications running on the remote system to display their windows on the local system.

upvoted 2 times

🗳️ 👤 **TheRealManish** 2 years, 1 month ago

Selected Answer: D

I pick D. the only hesitation I have is that installing X11 should set that automatically.

upvoted 2 times

🗳️ 👤 **e418137** 10 months, 3 weeks ago

It's a weird question these days. Imagine you have the workstation running X11. You use SSH to connect to a server, and you find software that you want to run. It's graphical software that won't run because it cannot find a display. On the remote server, you would 'export DISPLAY=hostname_or_ip_address:display_number' to cause the software to send its drawing commands to your local display.

upvoted 1 times

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.

Suggested Answer: BC

Community vote distribution

BC (77%)

AB (15%)

8%

🗳️ **Nvoid** Highly Voted 2 years, 1 month ago

Selected Answer: BC

B: Create unit file in /etc/systemd/system/

C: Add #!/bin/bash to the top of the script.

upvoted 6 times

🗳️ **dalv_01** Most Recent 8 months, 2 weeks ago

Selected Answer: A

in C is indicate wrong syntax "#!/bin/bash" double //

upvoted 1 times

🗳️ **wait4thebus** 1 year, 1 month ago

Selected Answer: BC

C is correct because every bash script should start with #!/bin/bash

B is correct because /etc/systemd/system is the correct folder for a unit file.

E is INCORRECT because /etc/init.d is an invalid folder path. E appears to be there to through people off the track of the correct answer.

upvoted 2 times

🗳️ **Alizadeh** 1 year, 4 months ago

Selected Answer: AB

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

upvoted 1 times

🗳️ **linux_admin** 1 year, 10 months ago

Selected Answer: BC

C. Add #!/bin/bash to the top of the script.

B. Create a unit file for the new service in /etc/systemd/system/ with the name helpme.service in the location.

Option C is necessary because it specifies the interpreter for the script. The shebang line, "#!/bin/bash", must be added to the top of the script to indicate which interpreter should be used to run the script.

Option B is necessary because the script must be registered as a service in the systemd system in order to be executed at startup. The administrator

should create a unit file for the new service in `/etc/systemd/system/` with the name `helpme.service`, which will describe how the service should be executed and how it should be managed by the system.

upvoted 2 times

🗨️ 👤 **Ckl22** 2 years ago

Selected Answer: BC

BC are the correct answers

upvoted 1 times

🗨️ 👤 **ryanzou** 2 years, 1 month ago

Selected Answer: AB

I think the answer is AB

upvoted 1 times

🗨️ 👤 **Veteran903** 2 years, 1 month ago

NO!, its B and C, `#!/bin/bash` is the very first line at the beginning on the script not the bottom

upvoted 8 times

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

Suggested Answer: *D*

🗨️ 👤 **linux_admin** 10 months, 2 weeks ago

D. chmod 744 filename.log

The chmod command is used to change the permissions of files in Linux. The number passed to chmod specifies the binary representation of the desired permissions. The binary representation of -rwxr--r-- is 111 100 100, which corresponds to the decimal value 744. So, the command to set the permissions to -rwxr--r-- is "chmod 744 filename.log".

upvoted 3 times

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

Suggested Answer: D

Community vote distribution

D (100%)

linux_admin 10 months, 2 weeks ago

D. chage -E -1 accountname

The chage command is used to modify the aging information of a user account, including the expiration date. The -E option is used to specify the expiration date of an account, and the -1 option sets the expiration date to a value of "never". So, the command "chage -E -1 accountname" will remove the expiration date of the account "accountname".

upvoted 3 times

[Removed] 10 months, 3 weeks ago

Selected Answer: D

-E is for expiration

upvoted 1 times

Ckl22 1 year ago

Selected Answer: D

D is the correct answer

upvoted 1 times

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

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

Suggested Answer: A

Community vote distribution

A (92%)

8%

🗳️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: A

The correct answer is A. visudo -c.

upvoted 3 times

🗳️ 👤 **linux_admin** 1 year, 4 months ago

Selected Answer: A

A. visudo -c

The visudo command is used to edit the /etc/sudoers file, and the -c option is used to check the syntax of the file without actually making any changes to it. If there are any syntax errors in the file, visudo will return an error message and prevent the changes from being saved. So, the command "visudo -c" can be used to validate the sudo rules in /etc/sudoers.

Option B "test -f /etc/sudoers" checks if the file /etc/sudoers exists, but it does not validate the syntax of the file.

Option C "sudo vi check" opens the file "check" with the vi editor, but it does not relate to the /etc/sudoers file.

Option D "cat /etc/sudoers | tee test" displays the contents of the /etc/sudoers file and saves it to a file named "test", but it does not validate the syntax of the file.

upvoted 4 times

🗳️ 👤 **linux_admin** 1 year, 4 months ago

sudo visudo -c

/etc/sudoers: parsed OK

/etc/sudoers.d/README: parsed OK

/etc/sudoers.d/kali-grant-root: parsed OK

upvoted 2 times

🗳️ 👤 **Ckl22** 1 year, 6 months ago

Selected Answer: C

C is the correct answer

upvoted 1 times

🗳️ 👤 **Ckl22** 1 year, 6 months ago

Aamm033 is right, its A.

upvoted 2 times

🗳️ 👤 **Aamm033** 1 year, 8 months ago

Selected Answer: A

-c flag is to check-only mode

upvoted 4 times

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

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

Suggested Answer: D

Community vote distribution

D (100%)


 **Ckl22** Highly Voted 1 year, 6 months ago

Selected Answer: D

D is the correct command but is missing the path to the public key that is to be copied over

```
$ ssh-copy-id -i ~/.ssh/id_rsa.pub user@host_address
```

upvoted 7 times

 **Alizadeh** Most Recent 10 months, 2 weeks ago

Selected Answer: D

The correct answer is D. `ssh-copy-id user@server`.
upvoted 1 times

A Linux administrator created a new file system. Which of the following files must be updated to ensure the filesystem mounts at boot time?

- A. /etc/sysctl
- B. /etc/filesystems
- C. /etc/fstab
- D. /etc/nfsmount.conf

Suggested Answer: C

  **linux_admin** 10 months, 2 weeks ago

C. /etc/fstab

The /etc/fstab file is used to specify the file systems that are mounted at boot time in Linux. The administrator must add an entry for the new file system to this file, which includes information such as the file system device, mount point, file system type, and mount options. This information is used by the mount command to mount the file system at boot time.

upvoted 4 times

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

```
$ vmstat -s --unit M
```

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

```
$ free -h
```

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

```
$ ps -aux | grep script.sh
```

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

Which of the following commands would address the issue?

- A. top -p 8321
- B. kill -9 8321
- C. renice -10 8321
- D. free 8321

Suggested Answer: C

Community vote distribution

B (91%) 9%

Towjumper Highly Voted 7 months, 3 weeks ago

Selected Answer: B

Question like this are why people despise CompTIA and their ridiculously worded questions.
upvoted 7 times

DRVision Most Recent 1 year ago

Selected Answer: B

obviously b, lowering the nice value from the default 0 and increasing the priority would dedicate more resources
upvoted 2 times

zionjr 7 months ago

A nice value of 10 indicates a process with lower priority than the default (0).
so it can actually free up some memory
upvoted 1 times

Damon54 1 year, 3 months ago

Selected Answer: B

No, corroge I think B is actually better as an answer
upvoted 1 times

Damon54 1 year, 3 months ago

Selected Answer: C

renice -10 lowers the priority of the process with PID 8321 by changing its nice value. A lower nice value means the process will have a higher priority, which can help prevent it from consuming too much CPU and memory resources. This command won't kill the process but will make it less aggressive in resource consumption.
upvoted 1 times

Zimendrakon 1 year, 2 months ago

A lower priority value means the process demands more resources, possibly denying those resources to processes that are "nicer".

upvoted 1 times

  **Zimendrakon** 1 year, 2 months ago

A Linux administrator is troubleshooting a memory-related issue- but what is the issue that the process needs more memory or that we need to free memory ?



upvoted 2 times

  **Jacobmy98** 1 year, 8 months ago

Selected Answer: B

b is it

upvoted 1 times

  **linux_admin** 1 year, 10 months ago

Selected Answer: B

B. kill -9 8321

If a process is causing memory issues, it may be necessary to terminate the process. The "kill" command is used to send signals to processes, and the -9 option sends the SIGKILL signal, which terminates the process immediately. By using the command "kill -9 8321", the administrator can terminate the process causing memory issues.

upvoted 3 times

  **MissAllen** 2 years, 1 month ago

Answer B is correct. The renice command speeds up or slows down processes in regards to CPU access. Killing the process would free up the memory.

upvoted 3 times

  **TheRealManish** 2 years, 1 month ago

this question is ridiculous. 1/2 of the total memory is listed as available.. there is not memory problem here.. but yes.. B is the only way to free up memory out of these answers..

upvoted 2 times

A systems administrator made some unapproved changes prior to leaving the company. The newly hired administrator has been tasked with revealing the system to a compliant state. Which of the following commands will list and remove the correspondent packages?

- A. dnf list and dnf remove last
- B. dnf remove and dnf check
- C. dnf info and dnf upgrade
- D. dnf history and dnf history undo last

Suggested Answer: D

Community vote distribution

D (100%)

🗲️ 👤 **MaryamNesa** Highly Voted 1 year, 1 month ago

D is correct. check the following link:

https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html/managing_software_with_the_dnf_tool/assembly_handling-package-manage-history_managing-software-with-the-dnf-

tool#:~:text=Reverting%20transactions%201%20To%20revert%20a%20particular%20transaction%2C,last%20transaction%2C%20use%3A%20%23%20dnf%20h
upvoted 6 times

🗲️ 👤 **linux_admin** Most Recent 10 months, 2 weeks ago

Selected Answer: D

D. dnf history and dnf history undo last

The "dnf history" command displays the history of all transactions performed using the DNF package manager, including package installs, upgrades, and removals. The "dnf history undo last" command undoes the last transaction, effectively removing the package that was installed. By using these two commands, the systems administrator can reveal the system to a compliant state by removing the unapproved packages.

Option A "dnf list and dnf remove last" lists the packages installed on the system, but it does not undo the last transaction. The "dnf remove last" command is not a valid command.

upvoted 4 times

🗲️ 👤 **Ckl22** 1 year ago

Selected Answer: D

Yeah D is the correct answer.

If there were multiple unwanted changes, the admin could also run the "rollback" option instead of "last"

upvoted 1 times

🗲️ 👤 **Veteran903** 1 year, 1 month ago

D is wrong, the right answer is A

upvoted 1 times

🗲️ 👤 **TheRealManish** 1 year, 1 month ago

I tested it on my VM and D worked for me.. as for a when do - dnf remove last, it says "no mach for argument last"

upvoted 2 times

🗲️ 👤 **Veteran903** 1 year, 1 month ago

you are correct, D is the right answer

upvoted 1 times

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

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

Which of the following commands would resolve the issue?

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

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ **HappyDay030303** 2 months, 1 week ago

Selected Answer: D

file permissions on `.ssh/authorized_keys` must be restricted for SSH to accept it
`-rw-r--r--` the last `r` here means readable by others
it is insecure so ssh ignores it
upvoted 1 times

🗳️ **e418137** 10 months, 3 weeks ago

Selected Answer: A

(A) is most plausibly correct. A typical SELinux label for `authorized_keys` looks like this: `'unconfined_u:object_r:ssh_home_t:s0'`. (B) There is no command, `'ssh_keygen'`. (It's `'ssh-keygen'`. (C) The file is owned by its user, not root. (D) The file, `authorized_keys`, works fine in 0600, 0640, and 0644. (But it's true that the best practice is 0600.)
upvoted 1 times

🗳️ **salrtom** 1 year, 2 months ago

D. `chmod 600 .ssh/authorized_keys`. It must have 600 permissions to work properly.
upvoted 2 times

🗳️ **LKRISB** 1 year, 7 months ago

D. `chmod 600 .ssh/authorized_keys`

This command sets the file permissions of `.ssh/authorized_keys` to 600, which means that only the owner (in this case, the user) will have read and write permissions, and no other users will have any permissions.

By setting the correct permissions on the `authorized_keys` file, SSH will be able to use the key for authentication without asking for a password.
upvoted 2 times

🗳️ **linux_admin** 1 year, 10 months ago

Selected Answer: A

"`restorecon`" is a command in SELinux (Security-Enhanced Linux) that is used to reset the security context of a file to its default SELinux security context. The "`restorecon`" command can be useful in cases where the SELinux security context of a file has been altered or changed, causing issues with the file's behavior or access.

The "`restorecon .ssh/authorized_keys`" command specifically resets the security context of the "`authorized_keys`" file in the ".ssh" directory to its default SELinux security context. This can be useful in cases where the SELinux security context of the "`authorized_keys`" file has been altered, causing issues with SSH authentication.
upvoted 4 times

🗳️ **Ckl22** 2 years ago

Selected Answer: A

A does appear to be the correct method to restore the default context of a file
upvoted 2 times

A cloud engineer needs to remove all dangling images and delete all the images that do not have an associated container. Which of the following commands will help to accomplish this task?

- A. docker images prune -a
- B. docker push images -a
- C. docker rmi -a images
- D. docker images rmi --all

Suggested Answer: A

Community vote distribution

A (100%)

linux_admin **Highly Voted** 10 months, 2 weeks ago

Selected Answer: A

A. docker images prune -a

The "docker images prune -a" command will remove all dangling images, which are images that are not associated with any existing containers, as well as all images that do not have an associated container. The "-a" option specifies that all images should be removed, regardless of whether they are in use or not.

upvoted 5 times

ckl22 **Most Recent** 1 year ago

Selected Answer: A

A is the correct answer:

\$ docker image prune [OPTIONS]

--all, -a Remove all unused images, not just dangling ones

upvoted 1 times

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

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

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

- A. Execute `grub-install --root-directory=/mnt` and reboot.
- B. Execute `grub-install /dev/sdX` and reboot.
- C. Interrupt the boot process in the GRUB menu and add `rescue` to the kernel line.
- D. Fix the partition modifying `/etc/default/grub` and reboot.
- E. Interrupt the boot process in the GRUB menu and add `single` to the kernel line.
- F. Boot the system on a LiveCD/ISO.

Suggested Answer: BD

Community vote distribution

BF (100%)

🗨️ **hackeriam1** 2 months, 1 week ago

Selected Answer: AF

So What's the "CompTIA-Correct" Answer?

According to CompTIA exam philosophy:

F. Boot system on a LiveCD/ISO → ✓ Absolutely correct and expected

A. `grub-install --root-directory=/mnt` and reboot → ✓ Also correct in the eyes of the exam

⚠️ Why Not B on the Exam?

`grub-install /dev/sdX` assumes a working mount or `chroot` – CompTIA wants the “safe, explicit method”.

They want to see you understand the full GRUB repair process:

Boot LiveCD

Mount system manually

Use `--root-directory=/mnt` to avoid `chrooting`

Even though B is technically valid, CompTIA tends to prefer explicit and less assumption-based commands.

upvoted 1 times

🗨️ **Rob74613** 7 months ago

Selected Answer: BF

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

This command reinstalls the GRUB bootloader on the specified disk (`/dev/sdX`). Replace `/dev/sdX` with the appropriate device identifier for the system's boot disk (e.g., `/dev/sda`, `/dev/nvme0n1`).

After executing the command, reboot the system and check if it successfully boots.

F. Boot the system on a LiveCD/ISO.

Booting the system from a LiveCD/ISO allows you to access the system's filesystem and repair the GRUB bootloader.

Once booted into the LiveCD/ISO environment, you can mount the system's root partition, chroot into it, and then reinstall GRUB using the appropriate commands (e.g., grub-install and update-grub).

After performing the necessary repairs, restart the system and check if it boots correctly.

upvoted 2 times

🗨️ 👤 **LKRISB** 7 months, 1 week ago

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

This action involves reinstalling the GRUB bootloader on the specified device (/dev/sdX). By reinstalling GRUB, it can help resolve any issues related to the bootloader and allow the system to boot successfully.

F. Boot the system on a LiveCD/ISO.

Booting the system using a LiveCD/ISO allows accessing the system's filesystem from an external environment. Once booted into the LiveCD/ISO, you can perform various troubleshooting steps such as checking the filesystem, repairing any errors, or modifying configuration files if necessary.

upvoted 1 times

🗨️ 👤 **linux_admin** 10 months, 2 weeks ago

Selected Answer: BF

Option F "Boot the system on a LiveCD/ISO" would help along with option B, as booting the system on a LiveCD/ISO would allow the administrator to access the file system and make necessary changes to fix the issue.

Additionally, the administrator could use the bootable LiveCD/ISO to run the grub-install command with the appropriate arguments, such as "/dev/sdX" (Option B), to install GRUB on the appropriate disk. This would help resolve the issue with the missing partition and allow the system to boot properly.

upvoted 4 times

🗨️ 👤 **KnifeClown1** 10 months, 3 weeks ago

Selected Answer: BF

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

F)Boot the system on a LiveCD/ISO.

upvoted 1 times

🗨️ 👤 **Ckl22** 1 year ago

Selected Answer: BF

I think its BF, but most resources state to:

1. Boot from live CD or ISO
2. Mount root partition
3. Reinstall GRUB

upvoted 1 times

🗨️ 👤 **TheRealManish** 1 year, 1 month ago

Does anyone have any idea if B and D are correct?

upvoted 2 times

A Linux administrator needs to create an image named sda.img from the sda disk and store it in the /tmp directory. Which of the following commands should be used to accomplish this task?

- A. `dd of=/dev/sda if=/tmp/sda.img`
- B. `dd if=/dev/sda of=/tmp/sda.img`
- C. `dd -if=/dev/sda --of=/tmp/sda.img`
- D. `dd --of=/dev/sda -if=/tmp/sda.img`

Suggested Answer: B

Community vote distribution

B (100%)

🗲️ 👤 **bongobo** 11 months ago

I-input O-output
upvoted 2 times

🗲️ 👤 **linux_admin** 1 year, 10 months ago

The correct command is B: `dd if=/dev/sda of=/tmp/sda.img`.

The dd command is used to create an image of a disk or partition, and the syntax is as follows: `dd if=input-file of=output-file`. In this case, the input file is the sda disk `/dev/sda` and the output file is the image file in the /tmp directory `/tmp/sda.img`.

upvoted 3 times

🗲️ 👤 **Ckl22** 2 years ago

Selected Answer: B

B is the only one that looks like the correct command syntax

upvoted 3 times

A Linux administrator is creating a primary partition on the replacement hard drive for an application server. Which of the following commands should the administrator issue to verify the device name of this partition?

- A. `sudo fdisk /dev/sda`
- B. `sudo fdisk -s /dev/sda`
- C. `sudo fdisk -l`
- D. `sudo fdisk -h`

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: C

The correct answer is C. `sudo fdisk -l`.
upvoted 1 times

🗳️ 👤 **linux_admin** 1 year, 4 months ago

The correct command is C. `sudo fdisk -l`

The `sudo fdisk -l` command is used to list information about all the partitions on the system. The output will include the device name, file system type, size, and other relevant information for each partition. In this case, the administrator is looking for the device name of the primary partition on the replacement hard drive for an application server. By using the `sudo fdisk -l` command, the administrator will be able to see a list of all the partitions on the system and identify the device name for the new partition.

upvoted 2 times

🗳️ 👤 **Ckl22** 1 year, 6 months ago

Selected Answer: C

"-l" = "--list"

List the partition tables for the specified devices and then
exit

upvoted 1 times

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

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

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

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

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

Which of the following is causing the issue?

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

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: A

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

upvoted 2 times

🗳️ 👤 **linux_admin** 1 year, 4 months ago

Selected Answer: A

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

```
#nameserver 10.10.10.10
```

upvoted 2 times

🗳️ 👤 **CKI22** 1 year, 6 months ago

Selected Answer: A

Whenever you want to comment a line, put a # in an appropriate place in a file. Anything beginning after # and ending at the end of the line won't get executed. This comments out the complete line. This comments out only the last part of the line starting at #

upvoted 2 times

A systems administrator is tasked with installing GRUB on the legacy MBR of the SATA hard drive. Which of the following commands will help the administrator accomplish this task?

- A. grub-install /dev/hda
- B. grub-install /dev/sda
- C. grub-install /dev/sr0
- D. grub-install /dev/hd0,0

Suggested Answer: B

Community vote distribution

B (100%)

linux_admin **Highly Voted** 10 months, 2 weeks ago

Selected Answer: B

The GRUB boot loader is a crucial component of a Linux system. It is responsible for loading the operating system into memory and allowing the user to select the desired operating system to boot into. In order to install GRUB on a legacy MBR (Master Boot Record) of a SATA hard drive, the administrator needs to specify the correct device name. The device name is used to identify the hard drive in the system, and it must be specified correctly in order for GRUB to be installed and function properly.

The correct device name to use in this scenario would be /dev/sda. The /dev/sda device name is the standard name used to identify the first SATA hard drive on a system. This is the device name that should be specified in the grub-install command

upvoted 5 times

abrilo **Most Recent** 1 year ago

Selected Answer: B

<https://www.thegeekdiary.com/grub-install-command-options/>

upvoted 2 times

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

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

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

- A. systemctl isolate multi-user.target
- B. systemctl isolate graphical.target
- C. systemctl isolate network.target
- D. systemctl isolate basic.target

Suggested Answer: B

🗨️ 👤 **Qubert2** 8 months, 2 weeks ago

he systemctl isolate command switches the system to the specified target. In this case, graphical.target starts all the services required for graphical mode (GUI) and stops services that are not needed in this mode.

It does not change the default runlevel, meaning the system will revert to its original default target upon the next reboot unless you change it using systemctl set-default.

upvoted 1 times

🗨️ 👤 **Jb4** 1 year, 4 months ago

Is it suppose to be A?

upvoted 3 times

🗨️ 👤 **angellorv** 2 years, 2 months ago

System V runlevel (Purpose)

runlevel 0 - (for System shutdown)

runlevel1 (for Single-user mode)

runlevel2 (for Local multiuser without remote network)

runlevel3 (for Full multiuser with network)

runlevel4 (for Unused/User-defined)

runlevel5 (for Full multiuser with network and display manager)

runlevel6 (for System reboot)

upvoted 4 times

🗨️ 👤 **angellorv** 2 years, 2 months ago

A runlevel is an operating state on a Unix and Unix-based operating system that is preset on the Linux-based system. Runlevels are numbered from zero to six.

Runlevels determine which programs can execute after the OS boots up. The runlevel defines the state of the machine after boot.

Systems administrators set the default runlevel of a system according to their needs, or use the runlevel command to find out the machine's current runlevel to assess a system. Runlevel 0 shuts down the system

Runlevel 1 single-user mode

Runlevel 2 multi-user mode without networking

Runlevel 3 multi-user mode with networking

Runlevel 4 user-definable

Runlevel 5 multi-user mode with networking

Runlevel 6 reboots the system to restart it

upvoted 1 times

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

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

Suggested Answer: A

Community vote distribution

A (100%)

🗨️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: A

The correct answer is A. `sudo useradd -e 2021-09-30 Project_user`.
upvoted 2 times

🗨️ 👤 **linux_admin** 1 year, 4 months ago

Selected Answer: A

The correct command to accomplish the task is A: `sudo useradd -e 2021-09-30 Project_user`. The -e option is used to specify the expiration date for the user account.
upvoted 2 times

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

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

Suggested Answer: A

Community vote distribution

A (100%)

🗨️ 👤 **linux_admin** 10 months, 2 weeks ago

Selected Answer: A

The correct command to download a Git repository from `https://git.company.com/admin/project.git` is:

A. `git clone https://git.company.com/admin/project.git`
upvoted 2 times

🗨️ 👤 **Ckl22** 1 year ago

Selected Answer: A

`git clone` is the best option if its a new project. `git pull` would work best if it was an ongoing project that had multiple team members that had pushed their updates to the main repository
upvoted 4 times

An administrator installed an application from source into `/opt/operations1/` and has received numerous reports that users are not able to access the application without having to use the full path `/opt/operations1/bin/*`. Which of the following commands should be used to resolve this issue?

- A. `echo 'export PATH=$PATH:/opt/operations1/bin' >> /etc/profile`
- B. `echo 'export PATH=/opt/operations1/bin' >> /etc/profile`
- C. `echo 'export PATH=$PATH/opt/operations1/bin' >> /etc/profile`
- D. `echo 'export $PATH:/opt/operations1/bin' >> /etc/profile`

Suggested Answer: A

Community vote distribution

A (100%)

🗨️ 👤 **linux_admin** 10 months, 2 weeks ago

Selected Answer: A

This command adds the path `/opt/operations1/bin` to the `$PATH` environment variable in the `/etc/profile` file. This will allow users to access the application without having to use the full path `/opt/operations1/bin/*`.

upvoted 4 times

🗨️ 👤 **Ckl22** 1 year ago

Selected Answer: A

A looks like the best answer here. It would be better to put the directory before `$PATH`, so that way the directory is searched before the rest of the pathing options in `$PATH`.

upvoted 1 times

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

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

Suggested Answer: DEH

Community vote distribution



Damon54 Highly Voted 1 year, 11 months ago

Selected Answer: DEH

`pvcreate /dev/sdb`
`vgextend /dev/rootvg /dev/sdb`
`lvresize -L +10G -r /dev/mapper/rootvg-rootlv`
 upvoted 9 times

IFBBPROSALCEDO Highly Voted 1 year, 1 month ago

Selected Answer: AEH

A. `df -h /`: To check the disk usage of the root filesystem.
 E. `lvresize -L +10G -r /dev/mapper/rootvg-rootlv`: To resize the logical volume and filesystem.
 H. `vgextend /dev/rootvg /dev/sdb`: To extend the volume group with a new physical volume if needed.
 upvoted 5 times

Misterymigi Most Recent 2 months, 3 weeks ago

Selected Answer: AEH

I would also like to think you would need to check first with `df -h /`
 then do `lvresize` along with `vgextend`

I don't see the point in creating a whole other physical volume when you can just resize the logical volume and volume group.

upvoted 1 times

Qubert2 8 months, 1 week ago

Selected Answer: DEH

Before you can extend the volume group with `vgextend`, you need to run `pvcreate` on the new disk/partition. Without this, the disk will not be recognized as part of LVM.
 upvoted 2 times

e418137 1 year, 4 months ago

Selected Answer: DEH

DEH. But in this order: DHE.
 D. '`pvcreate /dev/sdb`' (Create a physical volume for LVM on the new disk.)
 H. '`vgextend /dev/rootvg /dev/sdb`' (Extend volume group to include new disk.)
 E. '`lvresize -L +10G -r /dev/mapper/rootvg-rootlv`' (Add space to logical volume.)
 upvoted 2 times

sademik 1 year, 5 months ago

Selected Answer: AEH

Confirm, then resize.

upvoted 2 times

  **DRVision** 1 year, 6 months ago

Selected Answer: ACE

A. `df -h /`: This command will display the disk usage of the root filesystem in a human-readable format. It's a good starting point to understand how much space is being used.

C. `growpart /dev/mapper/rootvg-rootlv`: This command will resize the partition on the disk to use all available space. It's necessary if the filesystem is smaller than the underlying logical volume.

E. `lvresize -L +10G -r /dev/mapper/rootvg-rootlv`: This command will resize the logical volume to add an additional 10GB of space and resize the filesystem within the logical volume to use the additional space.

`fdisk -l /dev/sdb` is not a valid command, and `lsblk /dev/sda` would only list block devices and their sizes, but wouldn't help in resolving the issue.

Similarly, `parted -l /dev/mapper/rootvg-rootlv` would list partition layouts on a device, but wouldn't help in resolving a full filesystem. Finally, `vgextend /dev/rootvg /dev/sdb` and `pvcreeate /dev/sdb` would be used if you were adding a new physical disk to a volume group, which is not the case here.

upvoted 2 times

  **DRVision** 1 year, 6 months ago

CEH

A would only you to view the free space while H extends the volume group by adding partitions which would also reolve the issue.

upvoted 1 times

  **Damon54** 1 year, 11 months ago

Correct is

`pvcreeate /dev/sdb`

`vgextend /dev/rootvg /dev/sdb`

`lvresize -L +10G -r /dev/mapper/rootvg-rootlv`

upvoted 2 times

  **linux_admin** 2 years, 4 months ago

Selected Answer: AEH

A, E, and H are the commands that a systems administrator should use to resolve the issue of the root filesystem being full.

A. The `df -h /` command is used to check the disk usage and available space on the root filesystem.

E. The `lvresize -L +10G -r /dev/mapper/rootvg-rootlv` command is used to resize the logical volume to increase the root filesystem's size.

H. The `vgextend /dev/rootvg /dev/sdb` command is used to extend the root volume group with the new disk `/dev/sdb` in order to increase the root filesystem's size.

upvoted 3 times

  **linux_admin** 2 years, 4 months ago

Discard Im going with DEH.

upvoted 5 times

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

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

Suggested Answer: B

Community vote distribution

B (100%)

🗨️ 👤 **linux_admin** 10 months, 2 weeks ago

B. docker cp container_id:/deployment.yaml deployment.yaml
upvoted 2 times

🗨️ 👤 **Ckl22** 1 year ago

Selected Answer: B

Given answer appears correct(B)

The file is being copied from the container (SRC) to the host (DST)

docker cp [OPTIONS] CONTAINER:SRC_PATH DEST_PATH|-

docker cp [OPTIONS] SRC_PATH|- CONTAINER:DEST_PATH


upvoted 3 times

A Linux system is failing to start due to issues with several critical system processes. Which of the following options can be used to boot the system into the single user mode? (Choose two.)

- A. Execute the following command from the GRUB rescue shell: `mount -o remount, ro/sysroot`.
- B. Interrupt the boot process in the GRUB menu and add `systemd.unit=single` in the kernel line.
- C. Interrupt the boot process in the GRUB menu and add `systemd.unit=rescue.target` in the kernel line.
- D. Interrupt the boot process in the GRUB menu and add `single=user` in the kernel line.
- E. Interrupt the boot process in the GRUB menu and add `init=/bin/bash` in the kernel line.
- F. Interrupt the boot process in the GRUB menu and add `systemd.unit=single.target` in the kernel line.

Suggested Answer: CE

Community vote distribution



🗳️ 👤 **HappyDay030303** 2 months, 1 week ago

Selected Answer: BC

- B. `systemd.unit=single`
 - C. `systemd.unit=rescue.target`
- upvoted 2 times

🗳️ 👤 **NastyNutsu** 4 months, 1 week ago

Selected Answer: BC

- B. `systemd.unit=single`: This option boots the system directly into single user mode.
 - C. `systemd.unit=rescue.target`: This option boots the system into rescue mode, which is similar to single user mode and provides a minimal environment for troubleshooting.
- upvoted 2 times

🗳️ 👤 **Misterymigi** 4 months, 3 weeks ago

Selected Answer: BC

- I believe the question is just asking which of the two will boot into single user mode, in this case it will be B and C
- B- By adding `systemd.unit=single`, will tell the system to boot into single-user mode. It's another way of doing `rescue.target` which is also single-user mode.
 - C- By adding `systemd.unit=rescue.target`, this will put the system into rescue mode, which again is single-user mode

This is all based if the question is asking what it is asking which is "pick 2 which will boot the system into single user mode" nothing more but with compia's questioning, I think you never truly know what they want.

upvoted 1 times

🗳️ 👤 **Misterymigi** 2 months, 2 weeks ago

After careful review, im gonna go B and E

upvoted 1 times

🗳️ 👤 **[Removed]** 8 months ago

why isnt it BE?

upvoted 3 times

🗳️ 👤 **linux_admin** 10 months, 2 weeks ago

Selected Answer: DE

Option D (adding "`single=user`" in the kernel line) and option E (adding "`init=/bin/bash`" in the kernel line) can be used to boot the Linux system into single user mode. In single user mode, the system will only run a minimal set of services and the root user will be given a shell, allowing the administrator to troubleshoot and repair any issues with the system.

Option C, adding `systemd.unit=rescue.target` to the kernel line, is intended to boot the system into a rescue environment, not a single user mode. In

the rescue environment, the system will attempt to repair the system and restore the system to a bootable state, rather than providing direct access to the system as single user mode does.

upvoted 1 times

  **linux_admin** 10 months, 1 week ago

Discard I'm going with CE.

upvoted 2 times

  **Ckl22** 1 year ago

Selected Answer: CE

I think the given answer (CE) is correct

- C. "systems.unit=rescue.target" will put you in single-user mode

- E.

1. In GRUB, press E to edit your boot entry (the Ubuntu entry).

2. Look for the line that starts with linux, and then look for ro.

3. Replace ro with rw init=/bin/bash.

This action mounts your file system as read-write and uses /bin/bash as the init process.

4. Press Ctrl+X to reboot with these settings

upvoted 4 times


A DevOps engineer needs to allow incoming traffic to ports in the range of 4000 to 5000 on a Linux server. Which of the following commands will enforce this rule?

- A. `iptables -f filter -I INPUT -p tcp --dport 4000:5000 -A ACCEPT`
- B. `iptables -t filter -A INPUT -p tcp --dport 4000:5000 -j ACCEPT`
- C. `iptables filter -A INPUT -p tcp --dport 4000:5000 -D ACCEPT`
- D. `iptables filter -S INPUT -p tcp --dport 4000:5000 -A ACCEPT`

Suggested Answer: B

Community vote distribution

B (100%)

 **linux_admin** Highly Voted 10 months, 2 weeks ago

Selected Answer: B

The command `iptables` is used to manage the rules in the Linux kernel's firewall. The options used in the command determine how the rule will be enforced. In the case of option B, `-t filter` specifies that the rules should be applied to the filter table, which is used for packet filtering. The `-A INPUT` option specifies that the rule should be appended to the INPUT chain, which is used for incoming traffic. The `-p tcp` option specifies that the rule should only apply to TCP traffic, and the `--dport 4000:5000` option specifies that the rule should only apply to incoming traffic to ports in the range of 4000 to 5000. The `-j ACCEPT` option specifies that the matching traffic should be accepted, allowing it to enter the system.

upvoted 5 times

A Linux administrator needs to determine whether a hostname is in the DNS. Which of the following would supply the information that is needed?

- A. nslookup
- B. rsync
- C. netstat
- D. host

Suggested Answer: A

Community vote distribution

D (67%)

A (33%)

🗳️ 👤 **e418137** Highly Voted 10 months, 3 weeks ago

Both 'nslookup' and 'host' can resolve hostname, even against specific name servers. This is bad question.
upvoted 6 times

🗳️ 👤 **Monty97** Most Recent 9 months, 2 weeks ago

Selected Answer: D

Host is the newer version, nslookup is deprecated.
upvoted 2 times

🗳️ 👤 **sademik** 11 months ago

Selected Answer: A

nslookup
upvoted 1 times

🗳️ 👤 **Damon54** 1 year, 3 months ago

Selected Answer: D

nslookup - query Internet name servers (deprecated)
host - newer DNS lookup utility
upvoted 2 times

🗳️ 👤 **Damon54** 1 year, 4 months ago

host also works, what will be the correct answer ?
upvoted 2 times

🗳️ 👤 **linux_admin** 1 year, 10 months ago

Selected Answer: A

The correct command is "nslookup." The "nslookup" command allows a user to query a DNS server to determine the IP address of a hostname, or to determine the hostname associated with a particular IP address. It is commonly used to troubleshoot DNS resolution issues. The other options "rsync," "netstat," and "host" are not used for this purpose.
upvoted 1 times

🗳️ 👤 **e418137** 10 months, 1 week ago

From the manual: 'host - DNS lookup utility'
upvoted 1 times

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

Routing table:

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

IP configuration:

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

ARP table:

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

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


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

- A. An internal-only DNS server is configured.
- B. The IP netmask is wrong for ens3.
- C. Two default routes are configured.
- D. The ARP table contains incorrect entries.

Suggested Answer: C

Community vote distribution

B (100%)

 **Misteryyagi** 2 months, 3 weeks ago

Selected Answer: C

This is definitely C, you cant have 2 default gateways.

upvoted 1 times

 **Knocks** 3 months ago

Selected Answer: A

I would go with internal-only DNS: it is true that the gateway situation is weird (the gateway is outside the network, but reachable with an on-link static), but it is also true that the default gateway IS present in the ARP table (also with a complete flag), so the gateway should be working just fine. Likewise, having multiple default routes does add complexity, but also just works (and the routes have different distances).

We cannot say whether the ARP table contains incorrect entries or not.

The only remaining one is the DNS, which would explain why some websites work

upvoted 1 times

🗨️ 👤 **Damon54** 9 months, 3 weeks ago

Selected Answer: B

B. The IP netmask is wrong for ens3. - the IP of the GW is certainly wrong
C. Two default routes are configured. - actually having a double default can use problems
D. The ARP table contains incorrect entries. , you see two IP and MAC associated with the same network card...
upvoted 1 times

🗨️ 👤 **Ckl22** 1 year, 6 months ago

Selected Answer: B

Technically B is the best answer here, but its not best practice to have two default-routes configured. It overcomplicates the network configuration, and also the troubleshooting when an issue arises. A floating static route for a specific destination address is the better solution.
upvoted 1 times

🗨️ 👤 **Nvoid** 1 year, 7 months ago

oh shit!!! good call Manish!! i just realized that the real answer is D, looks at `ens3` interface in the arp table!! it's showing twice and it's the main default gateway, you're a genius!! i'm glad a looked at this again!
upvoted 2 times

🗨️ 👤 **TheRealManish** 1 year, 7 months ago

Selected Answer: B

well the 89 networks definitely have a subnet ID problem, but I'm not sure if that's the reason there is a problem here.
upvoted 4 times

🗨️ 👤 **Nvoid** 1 year, 7 months ago

Answer is D!!
upvoted 2 times

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

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

Suggested Answer: B

Community vote distribution

B (100%)

MissAllen **Highly Voted** 1 year, 7 months ago

B is best, but the IP address is wrong.

upvoted 9 times

Alizadeh **Most Recent** 10 months, 2 weeks ago

Selected Answer: B

The correct answer is B. `iptables -A INPUT -s 192.168.10.50 -j DROP`.

upvoted 1 times

linux_admin 1 year, 4 months ago

B. `iptables -A INPUT -s 192.168.10.30 -j DROP`

upvoted 1 times

Ckl22 1 year, 6 months ago

MissAllen is right, B is the answer here.

"-A" appends the rule, "-s" is source address, and "-j" is the action if the rule is matched, so DROP

upvoted 3 times

A Linux systems administrator is configuring a new filesystem that needs the capability to be mounted persistently across reboots. Which of the following commands will accomplish this task? (Choose two.)

- A. `df -h /data`
- B. `mkfs.ext4 /dev/sdc1`
- C. `fsck /dev/sdc1`
- D. `fdisk -l /dev/sdc1`
- E. `echo "/data /dev/sdc1 ext4 defaults 0 0" >> /etc/fstab`
- F. `echo "/dev/sdc1 /data ext4 defaults 0 0" >> /etc/fstab`

Suggested Answer: BF

Community vote distribution

BF (100%)

linux_admin 10 months, 2 weeks ago

Selected Answer: BF

`mkfs.ext4 /dev/sdc1` - This command creates the ext4 filesystem on the device `/dev/sdc1`.

`echo "/dev/sdc1 /data ext4 defaults 0 0" >> /etc/fstab` - This command appends a line to the `/etc/fstab` file that specifies the device to be mounted at a specific mount point persistently across reboots.

After running these commands, the new filesystem will be created and automatically mounted at the specified mount point every time the system reboots.

upvoted 3 times

Ckl22 1 year ago

Selected Answer: BF

Answer here is BF

"modify the `/etc/fstab` text file to automatically mount the new partition by opening it in an editor and adding the following line:

`/dev/ xxx 1 /data ext4 defaults 1 2`

where xxx is the device name of the storage device"

<https://learning.oreilly.com/library/view/mastering-linux-system/9781119794455/b01.xhtml>

upvoted 2 times

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

- A. parted
- B. df
- C. mount
- D. du
- E. fdisk
- F. dd
- G. ls

Suggested Answer: BD

Community vote distribution

BD (100%)

linux_admin 10 months, 2 weeks ago

Selected Answer: BD

B. df - This command shows the disk space usage of all mounted filesystems on the system. It will help identify the filesystems that are consuming the most space and the available space left.

D. du - This command shows the disk usage of files and directories, including hidden files. It can help identify which directories or files are using the most space on the filesystem.

upvoted 3 times

Ckl22 1 year ago

Selected Answer: BD

BD is the answer here.

du is a standard Unix program used to estimate file space usage

df is a standard Unix command used to display the amount of available disk space for file systems

upvoted 2 times

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

- A. fg
- B. su
- C. bg
- D. ed

Suggested Answer: A

Community vote distribution

A (100%)

linux_admin 10 months, 2 weeks ago

Selected Answer: A

The command the administrator can use to go back to the program is fg.

A. fg - This command brings a job to the foreground. In this case, the job is the program that was suspended using Ctrl+Z.

B. su - This command is used to switch to another user account.

C. bg - This command starts a suspended job in the background. It is not useful in this case because the administrator wants to bring the job to the foreground.

D. ed - This command starts the ed text editor. It is not useful in this case because the administrator wants to resume the program that was suspended.

upvoted 3 times

ckl22 1 year ago

Selected Answer: A

Answer is A. Ctrl+Z suspended the process, and "fg" will bring it back into the foreground of the shell

upvoted 1 times

A systems administrator received a notification that a system is performing slowly. When running the top command, the systems administrator can see the following values:

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

Which of the following commands will the administrator most likely run NEXT?

- A. vmstat
- B. strace
- C. htop
- D. lsof

Suggested Answer: A

Community vote distribution

A (100%)

🗨️ **linux_admin** 10 months, 2 weeks ago

Selected Answer: A

Based on the output of the top command, which shows a high percentage of CPU time spent on I/O wait (%wa), the command that the administrator would most likely run next is vmstat.

A. vmstat - This command provides a summary of system memory, CPU usage, and I/O statistics. It can be used to identify I/O bottlenecks and other performance issues.

B. strace - This command is used to trace system calls and signals made by a process. It is not directly relevant to identifying I/O bottlenecks.

C. htop - This command is an interactive process viewer that is similar to top but provides more features and capabilities. It may be useful for more detailed analysis of system performance, but is not necessary in this case.

D. lsof - This command lists open files and can be used to identify processes that are holding open file handles. It is not directly relevant to identifying I/O bottlenecks.

upvoted 4 times

🗨️ **Ckl22** 1 year ago

Selected Answer: A

A is the answer here, as the I/O wait time is 95%. vmstat will show the resource usage for each process

upvoted 1 times

Which of the following technologies provides load balancing, encryption, and observability in containerized environments?

- A. Virtual private network
- B. Sidecar pod
- C. Overlay network
- D. Service mesh

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: D

The answer is D. Service mesh.

upvoted 2 times

🗳️ 👤 **linux_admin** 1 year, 4 months ago

Selected Answer: D

A service mesh is a dedicated infrastructure layer that provides communication management between microservices in a containerized environment. Service meshes provide a range of features including load balancing, service discovery, traffic routing, encryption, and observability. The most popular service mesh frameworks are Istio, Linkerd, and Consul.

The other options provided are:

A. Virtual private network (VPN) - A VPN is a secure connection between two networks over the internet, which is primarily used to ensure secure remote access to a private network.

B. Sidecar pod - A sidecar pod is a design pattern in Kubernetes that allows for two containers to run within the same pod.

C. Overlay network - An overlay network is a virtual network that is built on top of an existing physical network, providing additional network management features, but does not inherently provide the features mentioned in the question.

upvoted 2 times

🗳️ 👤 **Ckl22** 1 year, 6 months ago

Selected Answer: D

Answer is D

"A service mesh controls the delivery of service requests in an application. Common features provided by a service mesh include service discovery, load balancing, encryption and failure recovery."

<https://www.techtarget.com/searchitoperations/definition/service-mesh>

upvoted 1 times

A development team asks an engineer to guarantee the persistency of journal log files across system reboots. Which of the following commands would accomplish this task?

- A. `grep -i auto /etc/systemd/journald.conf && systemctl restart systemd-journald.service`
- B. `cat /etc/systemd/journald.conf | awk '(print $1,$3)'`
- C. `sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i 'persistent/s/^#//q' /etc/systemd/journald.conf`
- D. `journalctl --list-boots && systemctl restart systemd-journald.service`

Suggested Answer: A

Community vote distribution

C (82%)

A (18%)

🗳️ 👤 tegami 11 months, 2 weeks ago

Selected Answer: C

The correct answer is C :)

upvoted 2 times

🗳️ 👤 Damon54 1 year, 2 months ago

Selected Answer: A

TEST KO `sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i 'persistent/s/^#//q' /etc/systemd/journald.conf`

TEST OK `sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i '/persistent/s/^#//' /etc/systemd/journald.conf`

Another method

TEST OK `sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i '/^#Storage/s/^#//' /etc/systemd/journald.conf`

upvoted 1 times

🗳️ 👤 Damon54 1 year, 3 months ago

Selected Answer: A

the correct syntax for the answer c should be

`sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i '/^#persistent/s/^#//' /etc/systemd/journald.conf`

At this point I would vote for answer A

upvoted 1 times

🗳️ 👤 Damon54 1 year, 4 months ago

Selected Answer: C

The supported values are "volatile", "persistent", "auto" and "none"

Default storage type is configured as "auto"

If "volatile", journal log data will be stored only in memory, i.e. below the /run/log/journal hierarchy (which is created if needed)

If "persistent", data will be stored preferably on disk, i.e. below the /var/log/journal hierarchy, with a fallback to /run/log/journal during early boot stage and if the disk is not writable

The "auto" value will configure journald to store journal log data in the /var/log/journal/ directory. However, the directory must already exist and have the proper permissions set. If it does not exist, then journal data is stored in the volatile /run/log/journal/ directory, and the data is erased when the system shuts down.

"none" turns off all storage, all log data received will be dropped.

upvoted 1 times

🗳️ 👤 LKRISB 1 year, 7 months ago

C. `sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i '/persistent/s/^#//' /etc/systemd/journald.conf`

Explanation:

The first part of the command, `sed -i 's/auto/persistent/g' /etc/systemd/journald.conf`, replaces all occurrences of "auto" with "persistent" in the /etc/systemd/journald.conf file. This ensures that the journal log files are set to be persistent.

The second part of the command, `sed -i '/persistent/s/^#//' /etc/systemd/journald.conf`, removes the comment character '#' from the line that contains "persistent" in the `/etc/systemd/journald.conf` file. This activates the persistence configuration.

By combining these two sed commands, the configuration file is modified to enable the persistence of journal log files across system reboots.

upvoted 4 times

🗨️ **nixonbii** 1 year, 10 months ago

Just asking a question, what guarantees that those changes in answer C will survive a system restart? Doesn't the file need to be saved and reloaded to become persistent?

upvoted 1 times

🗨️ **linux_admin** 1 year, 10 months ago

Selected Answer: C

The command `sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i 'persistent/s/^#//q' /etc/systemd/journald.conf` attempts to modify the `/etc/systemd/journald.conf` file to enable persistent journal storage.

Here is a breakdown of the command:

sed: a command-line utility for processing text

-i: a command-line option that instructs sed to edit the file in place, rather than outputting the edited text to the console

's/auto/persistent/g': the sed command to be executed. This command consists of three parts:

s: a command to substitute (replace) text

auto: the search pattern to look for in the file

persistent: the replacement pattern to replace the search pattern

g: a global option that means to replace all occurrences of the search pattern in the file, not just the first occurrence.

upvoted 2 times

🗨️ **linux_admin** 1 year, 10 months ago

The command `sed -i 'persistent/s/^#//q' /etc/systemd/journald.conf` attempts to remove a comment symbol # before the word "persistent" in the `/etc/systemd/journald.conf` file.

sed: a command-line utility for processing text

-i: a command-line option that instructs sed to edit the file in place, rather than outputting the edited text to the console

'persistent/s/^#//q': the sed command to be executed. This command consists of three parts:

persistent: the search pattern to look for in the file

/: a delimiter that separates the search pattern from the replacement pattern

^#//q: the replacement pattern and a command to quit after the first match is found. The ^# means "replace the first occurrence of a comment symbol at the beginning of the line with nothing," and the q means "quit immediately after the replacement is made." Note that the persistent search pattern and the # comment symbol are separated by a caret (^) character, which is a regular expression metacharacter that matches the start of a line.

upvoted 2 times

🗨️ **KnifeClown1** 1 year, 10 months ago

Selected Answer: C

C for real

upvoted 2 times

🗨️ **KnifeClown1** 1 year, 10 months ago

C. `sed -i 's/auto/persistent/g' /etc/systemd/journald.conf && sed -i 'persistent/s/^#//q' /etc/systemd/journald.conf`

upvoted 2 times

🗨️ **Pinnubhai** 1 year, 11 months ago

Selected Answer: C

Replace storage type with `Storage=persistent`

`# sed -i 's/#Storage.*/Storage=persistent/' /etc/systemd/journald.conf`

Next you can restart `systemd-journald` service

`# systemctl restart systemd-journald.service`

upvoted 1 times

🗨️ 👤 **Ckl22** 2 years ago

Selected Answer: C

C does appear to be the only one that would make the logs persistent

upvoted 1 times

🗨️ 👤 **MissAllen** 2 years, 1 month ago

I don't see any correct answer, but C is the closest. A is not correct because command one searches for the pattern "auto" in the file /etc/systemd/journald.conf, ignoring case. It will find it, thereby causing command two to execute. Restarting the journald service doesn't help if you made no changes to the .conf file!. Answer C searches for the pattern "auto", globally in the same file, but replaces it with the pattern "persistent", in the file (that is the -i option). That's great, but I tested it with command two and get an error. I suggest command two be the restart of the journald service.

upvoted 4 times

🗨️ 👤 **Nvoid** 2 years, 1 month ago

Good work MissAllen, (C) seems correct 1st command & (A) 2nd Command.

Heya wanna turn that Miss into Mrs.Allen? email: mrs.allen@blackhat.io <3 you can come live on this paradise island wit me :)

upvoted 4 times

A systems administrator is receiving tickets from users who cannot reach the application app that should be listening on port 9443/tcp on a Linux server.

To troubleshoot the issue, the systems administrator runs netstat and receives the following output:

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

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

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

Suggested Answer: B

Community vote distribution



B (100%)

  **tony12345** Highly Voted 9 months, 2 weeks ago

B because longest answer
upvoted 7 times

  **ajna_** 7 months, 3 weeks ago

lol...
upvoted 1 times

  **linux_admin** Most Recent 1 year, 4 months ago


Selected Answer: B

Based on the output of netstat, the issue is that the application is listening on the loopback interface (127.0.0.1) rather than on a publicly accessible IP address.

The netstat output shows that the application is bound to the IP address 127.0.0.1 on port 9443, which means that it is only accessible from the same system (localhost). The 0.0.0.0:* field in the output indicates that the application is listening on all available network interfaces, but the IP address 0.0.0.0 does not actually represent a valid IP address.

To resolve the issue, the application needs to be reconfigured to listen on a publicly accessible IP address, such as the IP address of the system's network interface that is reachable from the users' network.

upvoted 3 times

  **Ckl22** 1 year, 6 months ago

Selected Answer: B

Answer is B.

The server is in a "Listen" state on port 9943 using its loopback address. The "1234" is a process-id

upvoted 2 times

A systems administrator is troubleshooting a connectivity issue pertaining to access to a system named db.example.com. The system IP address should be 192.168.20.88. The administrator issues the dig command and receives the following output:

```
;; ANSWER SECTION:
db.example.com.      15 IN A 192.168.20.89
```

The administrator runs `grep db.example.com /etc/hosts` and receives the following output:

```
192.168.20.89 db.example.com
```

Given this scenario, which of the following should the administrator do to address this issue?

- A. Modify the `/etc/hosts` file and change the db.example.com entry to 192.168.20.89.
- B. Modify the `/etc/network` file and change the db.example.com entry to 192.168.20.88.
- C. Modify the `/etc/network` file and change the db.example.com entry to 192.168.20.89.
- D. Modify the `/etc/hosts` file and change the db.example.com entry to 192.168.20.88.

Suggested Answer: A

Community vote distribution

D (100%)

🗳️ **ryan zou** Highly Voted 2 years, 8 months ago

Selected Answer: D

Why not D?

upvoted 9 times

🗳️ **Veteran903** 2 years, 7 months ago

D is the right answer

upvoted 8 times

🗳️ **Nvoid** Highly Voted 2 years, 7 months ago

Selected Answer: D

D is correct.

upvoted 5 times

🗳️ **bc1235813** Most Recent 1 year, 3 months ago

None of the solutions provide a fix for the fact that remote systems are seeing the wrong ip address. The `/etc/hosts` file has no bearing on remote access. There is nothing that will fix this locally. The other systems need their `/etc/hosts` files fixed or DNS needs to be fixed.

upvoted 1 times

🗳️ **97155f3** 7 months, 1 week ago

the better answer is to remove everything from `/etc/hosts`, except the local hosts info and rely on dns. static entries for name resolution aren't optimal.

upvoted 1 times

🗳️ **97155f3** 7 months, 1 week ago

`cat /etc/nsswitch.conf`

hosts: files dns

upvoted 1 times

🗳️ **KnifeClown1** 2 years, 4 months ago

Selected Answer: D

D obviously

upvoted 2 times

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

Output 1:

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

Output 2:

```
nameserver 192.168.168.53
```

Output 3:

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

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

Output 4:

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

Output 5:

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

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

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

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

- A. The name www.comptia.org does not point to a valid IP address.
- B. The server 192.168.168.53 is unreachable.
- C. No default route is set on the server.
- D. The network interface eth0 is disconnected.

Suggested Answer: C

Community vote distribution

C (100%)

  **dsmitd33** 7 months, 2 weeks ago



Can someone please explain this better?

upvoted 2 times

  **ID77** 4 months ago

A default route is needed for a system to communicate with external networks (outside it's local subnet). There is no default route in output 4. It shows only a route for the local subnet 192.168.168.0/24. Without a default route, the system does not know where to send packets meant for external networks, which explains why it cannot reach www.comptia.org.

upvoted 1 times



  **Ckl22** 2 years, 6 months ago

Selected Answer: C

The answer should be C.

I can't see any default routes on this configuration. But the name server on the same subnet is reachable, and does appear to have a valid IP address. It looks like there is another typo on this question though. The broadcast address 192.168.169.255 is outside the subnet of 192.168.168.10/24. So any broadcasts at that IP would be useless when attempting to talk to 192.168.168.0/24 devices.

upvoted 4 times

  **Nvoid** 2 years, 7 months ago

Selected Answer: C

going with C.



upvoted 1 times

  **TheRealManish** 2 years, 7 months ago

Selected Answer: C

I'm picking C because i don't see any default route. I have no idea if the static name resolution is correct or not..

upvoted 1 times

  **Nvoid** 2 years, 7 months ago

process of elimination, C is left.

upvoted 3 times

A systems technician is working on deploying several microservices to various RPM-based systems, some of which could run up to two hours. Which of the following commands will allow the technician to execute those services and continue deploying other microservices within the same terminal section?

- A. gedit & disown
- B. kill %1
- C. fg %1
- D. bg %1 job name

Suggested Answer: A

Community vote distribution

D (91%)

9%

🗳️ 👤 **angellorv** 8 months, 3 weeks ago

Answer A:

bg puts existing running processes in the background; while "& disown" sends the process to the bg upon execution which is what the question is implying.

"...execute those services and continue deploying other microservices within the same terminal section?"

gedit - has nothing to do with the question or answer it's just the command used in the answer (bad choice); imagine answer A as: job name & disown
upvoted 3 times

🗳️ 👤 **linux_admin** 10 months, 2 weeks ago

Selected Answer: D

D. bg %1 job name: This command puts the job associated with job ID %1 in the background and allows the technician to continue deploying other microservices in the same terminal session. The job name argument is optional and allows the technician to specify a custom name for the job. This option is suitable for the scenario described in the question.

A. gedit & disown: This command opens the gedit text editor in the background and disowns it so that it is no longer associated with the terminal. This option is not relevant to deploying microservices and would not allow the technician to continue deploying other microservices in the same terminal session.

upvoted 2 times

🗳️ 👤 **KnifeClown1** 10 months, 3 weeks ago

Selected Answer: D

D. bg %1 job name

The command bg %1 job name will allow the technician to execute the microservice in the background, allowing them to continue deploying other microservices within the same terminal session. The bg command stands for "background" and it is used to run a job in the background. The %1 is a job specifier, which specifies the job to be run in the background, and "job name" is the name given to the job. This allows the technician to keep the microservice running even after they have closed the terminal session.

upvoted 1 times

🗳️ 👤 **Pinnubhai** 11 months, 3 weeks ago

Selected Answer: D

Answer D is the closest. A doesn't make any sense as what gedit will help in all this.

upvoted 4 times

🗳️ 👤 **Ckl22** 1 year ago

Selected Answer: D

D is the only one that makes the most sense. "bg %n" where n is the job ID, will send that job to process in the background, and all interaction with the current shell.

upvoted 2 times

🗳️ 👤 **lo_01234_ol** 1 year ago

Selected Answer: A

Scroll down to #5:



<https://www.baeldung.com/linux/detach-process-from-terminal>

upvoted 1 times

  **TheRealManish** 1 year ago

disown is fine, but what does gedit have to do with anything?


upvoted 3 times

  **Nvoid** 1 year, 1 month ago

Selected Answer: D

put the jobs in the background process!

upvoted 1 times

  **TheRealManish** 1 year, 1 month ago



it's hard to say..

bg%1 job name - is not a valid command.. i meant it works, but thats because the command ignores the job name.

also, the question says there are multiple jobs and the %1 would just background one of the jobs.



But then, B and C are obviously wrong. as for A, what the hell does gedit have to do with anything? ugh.. i guess D it is?

upvoted 1 times

  **Nvoid** 1 year, 1 month ago

best answer out of the worst of all possibilities. I'll come back around leave a comment after i pass.

upvoted 1 times

  **Nvoid** 1 year, 1 month ago

i came back around, but didn't take the test yet, i get what you're saying now about:

bg%1

and

bg%1 job_name

being the same command, hard to say..

nothing changed for me, terrible question or mis-worded?

upvoted 1 times

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

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

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

root@linux:~# free -m
              total    used     free   shared  buff/cache   available
Mem:           3933    1454       110       33       2368       2202
Swap:           1497         5       1491
```

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

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

Suggested Answer: B

Community vote distribution

B (70%)

D (30%)

 **DRVision** Highly Voted 1 year, 6 months ago

Selected Answer: B

Based on the output provided, the issue causing the slowness is B. The CPU is overloaded. The load average of 9.90 indicates that there are, on average, 9.90 processes in the running or runnable state. Given that there is only a single CPU, this means that there are far more processes needing CPU time than the CPU can handle, leading to an overloaded CPU. This is likely the cause of the slowness. The other options (A, C, and D) do not seem to be supported by the provided output. The output does not indicate that the system is running out of swap space (almost all free), that the memory is exhausted (still have a lot of free mem), or that the processes are paging (only 5 for pages used). Therefore, these are not likely to be the cause of the slowness.

upvoted 5 times


 **HappyDay030303** Most Recent 2 months, 1 week ago

Selected Answer: D

processes are paging

so (swap out) values 26211 and 5242 indicate active paging

upvoted 1 times

 **044f354** 9 months, 2 weeks ago

Selected Answer: D

D. The PROCESSES are PAGING.

The significant (>1000 KB/s) si (swap in from disk) and so (swap out to disk) values indicate that the system is actively PAGING data between RAM and swap space, which significantly slows down performance due to the slow nature of disk I/O compared to RAM access.

High values (>0 consistently) in the b column (blocked [aka waiting] PROCESSES) also support the indication of I/O bottlenecks related to PAGING.

Heavy PAGING indicates that PROCESSES are waiting for data to be read from or written to swap space, which causes the observed I/O bottleneck and system slowness. The system's high load averages and blocked PROCESSES further confirm this.

(caps for impact)

upvoted 1 times

🗨️ 👤 **Damon54** 1 year, 9 months ago

Selected Answer: D

While the load average is high, it doesn't directly imply that the CPU is overloaded. Instead, the "si" and "so" values in the vmstat output indicate that processes are paging in and out of swap, which typically points to I/O bottlenecks and disk activity as the primary cause of slowness.

So, based on the provided information, the most likely cause of slowness is related to I/O and processes paging, rather than CPU overload. Therefore, option B (The CPU is overloaded) is less likely to be the primary cause.

upvoted 1 times

🗨️ 👤 **Damon54** 1 year, 10 months ago

Selected Answer: D

B or D ? please help

As you can see we have bo values -> blocks sent to a block device. paged memory , and so value -> Amount of memory swapped to disk (/s).

and id -> idle elapsed time. = 0 (normal value is 99 and 100)

upvoted 1 times

🗨️ 👤 **Damon54** 1 year, 10 months ago

B definitive

id – The percentage of idle CPU.

upvoted 1 times

🗨️ 👤 **Landoski** 2 years, 5 months ago

The %id displays CPU idle time and if this is too high then this indicates that the CPU is working too hard

upvoted 1 times

🗨️ 👤 **TheRealManish** 2 years, 7 months ago

Selected Answer: B

im not seeing and swap or memory issues, what we do see is the CPU ID value at 99, so i pick B.

upvoted 4 times

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

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

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

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

Suggested Answer: D

Community vote distribution

D (78%)

B (22%)

🗳️ **dsmitd33** 7 months, 1 week ago

Wouldn't C also keep the permissions the same for the IT group while still giving the same permissions to the Finance group as D?
upvoted 1 times

🗳️ **bongobo** 1 year, 5 months ago

Debian12 / -bash: setfacl: command not found
upvoted 1 times

🗳️ **Alizadeh** 1 year, 10 months ago

Selected Answer: D

The correct answer is D. `setfacl -m g:finance:rw file`.
upvoted 1 times

🗳️ **linux_admin** 2 years, 4 months ago

Selected Answer: D

D. `setfacl -m g:finance:rw file`: This command sets a new ACL (access control list) on the file that grants read and write access to the finance group. This option would allow the finance department to access the file while also maintaining the existing permissions for IT employees. The + in the rw option means to add read and write permissions to the existing permissions. The g: specifies that the permission is being set for a group, and the finance is the name of the group.
upvoted 3 times

🗳️ **KnifeClown1** 2 years, 4 months ago

Selected Answer: D

The correct command to fix the access issue while limiting access to IT and finance department employees would be:

D. `setfacl -m g:finance:rw file`

This command sets a Access Control List (ACL) for the file with the permission "rw" (read and write) for the group "finance". This would allow the finance department employees to access the file with the necessary permissions, while the IT employees still retain their read and write permissions.
upvoted 2 times

🗳️ **TheRealManish** 2 years, 7 months ago

Selected Answer: D

if we change the owner to finance, the IT group would then be considered "other" in permissions and have no rights. currently IT has rights,, so using setfacl to give finance RW rights would keep the IT groups, and give finance rights. D
upvoted 2 times

🗳️ **Nvoid** 2 years, 7 months ago

good point, i wasn't sure why which would be the right answer, but this actual makes sense.

choosing D. here.

upvoted 1 times

  **lo_01234_ol** 2 years, 7 months ago

Selected Answer: B

chown [owner]:[group] file

This would allow rw for both IT and Finance.

upvoted 2 times

  **lo_01234_ol** 2 years, 6 months ago

Changing this to 'D.'

upvoted 1 times

  **TheRealManish** 2 years, 7 months ago

I disagree. if we change the owner to finance, the IT group would then be considered "other" in permissions and have no rights. currently IT has rights,, so using setfacl to give finance RW rights would keep the IT groups, and give finance rights. D

upvoted 1 times

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

- A. Create a unit file in the `/etc/default/` directory.
`systemctl enable cleanup`
`systemctl is-enabled cleanup`
- B. Create a unit file in the `/etc/ske1/` directory.
`systemctl enable cleanup`
`systemctl is-enabled cleanup`
- C. Create a unit file in the `/etc/systemd/system/` directory.
`systemctl enable cleanup`
`systemctl is-enabled cleanup`
- D. Create a unit file in the `/etc/sysctl.d/` directory.
`systemctl enable cleanup`
`systemctl is-enabled cleanup`

Suggested Answer: C

Community vote distribution

C (100%)

👤 **linux_admin** Highly Voted 1 year, 4 months ago

Selected Answer: C

To create a custom script, `cleanup.sh`, to run at boot as part of the system services, the process that would accomplish this task is option C, "Create a unit file in the `/etc/systemd/system/` directory".

Here's how this can be done:

Create the `cleanup.sh` script and place it in a suitable location. For example, in the `/usr/local/bin` directory.

Create a new systemd unit file with the `.service` extension in the `/etc/systemd/system` directory. For example, create a file called `cleanup.service`.

In the `cleanup.service` file, define the service by setting the service name, description, and the command to execute the script. For example:

[Unit]

Description=Cleanup script

[Service]

ExecStart=/usr/local/bin/cleanup.sh

[Install]

WantedBy=multi-user.target

upvoted 5 times

👤 **linux_admin** 1 year, 4 months ago

Save and close the `cleanup.service` file.

Use the `systemctl` command to enable the service to run at boot time:

`systemctl enable cleanup.service`

This will create the necessary symbolic links to start the `cleanup.service` unit automatically at boot time.

Use the `systemctl` command to check the status of the service:

systemctl is-enabled cleanup.service

This command will show whether the service is enabled to start at boot time or not.

After completing these steps, the cleanup.sh script will run automatically at boot time as part of the system services.
upvoted 2 times

  **Alizadeh** Most Recent 10 months, 2 weeks ago

Selected Answer: C

The correct answer is C. Create a unit file in the /etc/systemd/system/ directory.

upvoted 2 times

A Linux system is failing to boot. The following error is displayed in the serial console:

```
[[1;33mDEPEND[Om] Dependency failed for /data.
```

```
[[1;33mDEPEND[Om] Dependency failed for Local File Systems
```

...

Welcome to emergency mode! After logging in, type "journalctl -xb" to view system logs, "systemctl reboot" to reboot, "systemctl default" to try again to boot into default mode.

Give root password for maintenance

(or type Control-D to continue)

Which of the following files will need to be modified for this server to be able to boot again?

- A. /etc/mtab
- B. /dev/sda
- C. /etc/fstab
- D. /etc/grub.conf

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: C

The correct answer is C. /etc/fstab.

upvoted 1 times

🗳️ 👤 **nixonbii** 1 year, 4 months ago

This does not look like output that can be properly analyzed. Should we just assume what it says?:

```
[[1;33mDEPEND[Om] Dependency failed for /data.
```

```
[[1;33mDEPEND[Om] Dependency failed for Local File Systems
```

upvoted 1 times

🗳️ 👤 **linux_admin** 1 year, 4 months ago

Selected Answer: C

Based on the error message and symptoms described, the file that needs to be modified for this server to be able to boot again is option C, /etc/fstab.

The error message indicates that the system is failing to start the /data partition, which is causing the system to go into emergency mode and preventing it from booting. The /data partition is most likely defined in the /etc/fstab file, which contains information about the file systems that are mounted at boot time.

upvoted 4 times

🗳️ 👤 **KnifeClown1** 1 year, 4 months ago

Selected Answer: C

The file /etc/fstab is used by the Linux system to mount file systems during boot. It contains information about file systems to be mounted, including the device name, mount point, file system type, and mount options. The error message indicates that a dependency has failed for "/data" and for local file systems, which suggests that there is an issue with one of the file systems listed in the /etc/fstab file.

It is likely that the device name, mount point, file system type, or mount options for "/data" or another file system listed in /etc/fstab is incorrect or outdated, causing the system to fail to mount the file system during boot. The system administrator will need to access the /etc/fstab file and modify it to correct the issue and allow the system to boot successfully.

upvoted 2 times

🗳️ 👤 **lo_01234_ol** 1 year, 7 months ago

Selected Answer: C

<https://clay-atlas.com/us/blog/2021/07/25/linux-en-welcome-to-emergency-mode/>

upvoted 1 times

  **TheRealManish** 1 year, 7 months ago

Selected Answer: C

I pick C: fstab.. but i'm not really sure. . anyone else more sure?


upvoted 2 times

A systems administrator frequently connects to a remote host via SSH and a non-standard port. The systems administrator would like to avoid passing the port parameter on the command line every time. Which of the following files can be used to set a different port value for that host?

- A. /etc/ssh/sshd_config
- B. /etc/ssh/moduli
- C. ~/.ssh/config
- D. ~/.ssh/authorized_keys

Suggested Answer: C

Community vote distribution



linux_admin Highly Voted 1 year, 10 months ago

Selected Answer: C

The file that can be used to set a different port value for a specific remote host in SSH is option C, ~/.ssh/config.

The ~/.ssh/config file is a per-user configuration file that SSH reads before connecting to a remote host. It allows users to set various SSH options for specific hosts, eliminating the need to pass these options on the command line every time.

upvoted 5 times

e418137 Most Recent 10 months, 3 weeks ago

Selected Answer: A

Only the server side can change the port value in '/etc/ssh/sshd_config'.

upvoted 1 times

e418137 10 months, 3 weeks ago

Ugh... these questions! "Which of the following files can be used to set a different port value for that host?" Only the server side can change the port value in '/etc/ssh/sshd_config'. By specifying the port on the client side in '~/.ssh/config', one need not specify the port every time from the command line, but this does not "set a different port value for that host." SMH. I surely hope the test is not like this.

upvoted 2 times

Alizadeh 1 year, 4 months ago

Selected Answer: C

The correct answer is C. ~/.ssh/config.

upvoted 2 times

Ckl22 2 years ago

Selected Answer: C

Answer is C

For an individual user's connections to a remote system, create and/or modify the client side's ~/.ssh/config file.

For every user's connection to a remote system, create and modify the client side's /etc/ssh/ssh_config file.

upvoted 1 times

lo_01234_ol 2 years ago

Selected Answer: C

Bottom of page here:

<https://askubuntu.com/questions/1110326/how-to-connect-to-a-certain-port-on-ssh>

upvoted 2 times

MaryamNesa 2 years, 1 month ago



A is correct

upvoted 2 times

TheRealManish 2 years, 1 month ago

A is where the server side config resides. C is where the client side config resides. they are asking about changes to client, which happen on C

upvoted 4 times

  **Veteran903** 2 years, 1 month ago

Exactly!!

upvoted 3 times

A Linux administrator modified the SSH configuration file. Which of the following commands should be used to apply the configuration changes?

- A. systemctl stop sshd
- B. systemctl mask sshd
- C. systemctl reload sshd
- D. systemctl start sshd

Suggested Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: C

The correct answer is C. systemctl reload sshd.

upvoted 1 times

🗨️ 👤 **linux_admin** 1 year, 4 months ago

Selected Answer: C

The command that should be used to apply changes to the SSH configuration file is option C, systemctl reload sshd.

The reload command tells the sshd daemon to re-read its configuration file without terminating any existing connections. This makes it possible to apply changes to the SSH configuration file without disrupting any active SSH sessions.

Using the systemctl reload sshd command will cause the SSH service to reload its configuration file and apply any changes that have been made, without requiring the service to be stopped and restarted. This is the preferred method of applying changes to the SSH configuration file because it allows changes to be made without affecting any currently connected users.

If the systemctl reload sshd command is not available, the systemctl restart sshd command can be used instead. This will stop and start the SSH service, terminating all existing connections and applying the new configuration settings. However, this approach is less desirable because it will cause any active SSH sessions to be terminated.

upvoted 1 times

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

- A. ifconfig hw eth1
- B. netstat -r eth1
- C. ss -ti eth1
- D. ip link show eth1

Suggested Answer: D

Community vote distribution

D (75%)

B (25%)

🗳️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: D

The correct answer is D. ip link show eth1.

upvoted 1 times

🗳️ 👤 **angellorv** 1 year, 2 months ago

Deprecated Linux networking commands and their replacements: ip, netstat and others

<https://dougvitale.wordpress.com/2011/12/21/deprecated-linux-networking-commands-and-their-replacements/>

upvoted 1 times

🗳️ 👤 **linux_admin** 1 year, 4 months ago

Selected Answer: D

The command that can be used to check the link status of a network interface named eth1 in a Linux server is option D, ip link show eth1.

The ip command is used to show or manipulate the network interfaces on a Linux system. The link option is used to display or modify the attributes of a network device, and the show option is used to display the specified device.

To check the link status of eth1, the following command can be used:

ip link show eth1

This command will display information about the eth1 network interface, including its status and any configured options.

upvoted 1 times

🗳️ 👤 **KnifeClown1** 1 year, 4 months ago

Selected Answer: D

The command "ip link show eth1" is the correct command to check the link status of the network interface named eth1 in a Linux server.

The "ip" command is part of the iproute2 suite of networking tools, which provides a more advanced and flexible way of managing network interfaces compared to the traditional ifconfig command. The "link" subcommand of "ip" can be used to show information about network interfaces, including the link status. By adding the interface name (eth1), the command will only show information about the specified interface.

The "ip link show eth1" command will display information about the eth1 interface, including the current status of the link, the speed and duplex of the connection, and the MTU (maximum transmission unit) of the interface. If the link is up and running, the status of the interface will be "UP", and if there is an issue with the link, the status will be "DOWN".

upvoted 1 times

🗳️ 👤 **Ckl22** 1 year, 6 months ago

Selected Answer: B

Answer is B

netstat

upvoted 1 times

🗳️ 👤 **Ckl22** 1 year, 6 months ago

looks like ip link show is the correct command as netstat -r is meant to view route information, and not the status of a network interface

upvoted 3 times

A systems administrator is tasked with setting up key-based SSH authentication. In which of the following locations should the administrator place the public keys for the server?

- A. ~/.sshd/authkeys
- B. ~/.ssh/keys
- C. ~/.ssh/authorized_keys
- D. ~/.ssh/keyauth

Suggested Answer: C

Community vote distribution

C (100%)

🗲️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: C

The correct answer is C. ~/.ssh/authorized_keys.
upvoted 1 times

🗲️ 👤 **linux_admin** 1 year, 4 months ago

Selected Answer: C

The location where the systems administrator should place the public keys for the server to set up key-based SSH authentication is option C, ~/.ssh/authorized_keys.

The authorized_keys file is located in the .ssh directory in the home directory of the user account that is being used to log in to the remote server. This file contains a list of public keys that are allowed to authenticate the user when logging in to the server via SSH.

To set up key-based SSH authentication, the systems administrator should copy the public key(s) of the user(s) to the remote server's authorized_keys file using a secure method such as scp. Once the public key(s) have been added to the file, the user(s) should be able to log in to the server using their corresponding private key(s) without being prompted for a password.

upvoted 2 times

🗲️ 👤 **Ckl22** 1 year, 6 months ago

Selected Answer: C

Answer is C

"After that, copy the new public key to the server system's ~/.ssh/authorized_keys filesystem using the ssh-copy-id utility."


upvoted 1 times

A Linux administrator needs to create a new user named user02. However, user02 must be in a different home directory, which is under /comptia/projects. Which of the following commands will accomplish this task?

- A. useradd -d /comptia/projects user02
- B. useradd -m /comptia/projects user02
- C. useradd -b /comptia/projects user02
- D. useradd -s /comptia/projects user02

Suggested Answer: A

Community vote distribution



linux_admin 1 year, 4 months ago

Selected Answer: A

The command that can be used to create a new user named user02 with a different home directory under /comptia/projects is option A, useradd -d /comptia/projects user02.

The -d option is used to specify the home directory of the new user. By default, the useradd command creates the home directory of the new user in /home. To specify a different home directory for the new user, the -d option should be used, followed by the path to the new home directory.

upvoted 3 times

KnifeClown1 1 year, 4 months ago

Selected Answer: A

A. useradd -d /comptia/projects user02

The command useradd -d /comptia/projects user02 will create a new user named user02 and set its home directory to /comptia/projects/user02. The -d option is used to specify the home directory for the user. The useradd command is used to add a new user to the system, and it creates the user's home directory by default, so there is no need to use the -m option in this case.

upvoted 2 times

Ckl22 1 year, 6 months ago

Selected Answer: A

Answer is A

-d, --home HOME_DIR

The new user is created using HOME_DIR as the value for the user's login directory. The default is to append the LOGIN name to BASE_DIR and use that as the login directory name. The directory HOME_DIR does not have to exist but is not created if it's missing.

-b, --base-dir BASE_DIR

<https://www.computerhope.com/unix/useradd.htm>

upvoted 4 times

tony12345 1 year, 6 months ago

D is the only correct answer!!!! Noobz. Git good

upvoted 1 times

Notnotataco 1 year, 7 months ago

<https://unix.stackexchange.com/questions/83930/difference-between-useradd-b-and-useradd-d>

I really think the way this question is worded, the correct answer is C. How I'm reading it, it implies that /CompTIA/projects already exists and we just need to add this new user to that directory instead of the /home directory. The -d option creates a new directory, the -b option creates a directory for the user at the specified target

upvoted 1 times

Mathew89 6 months, 3 weeks ago

I think I'm going to have to agree with C as well. The -d flag sets the explicit directory so in this case that users home would be /comptia/projects
NOT /comptia/projects/user02

upvoted 1 times

  **Veteran903** 1 year, 6 months ago

Im with you.....C



upvoted 1 times

  **TheRealManish** 1 year, 7 months ago

Selected Answer: C

the wording on this one is really a problem. -d will store the users contents directly in the /compta/projects folder. the -b will create a comptia/projects/user2 folder and put it there. by the wording using the word UNDER, and sort of so simple logic, it seems like we should do C and have the -b option create /comptia/projects/user2

upvoted 1 times

  **Nvoid** 1 year, 7 months ago

naw have a look at this:

<https://linux.die.net/man/8/useradd>

if you specify -b the directory must exist or you must supply -d as well.

I went with A or B

-d won't create the home directory but -m will, so i think it's -m but for some reason i want to go with A.

upvoted 2 times

  **TheRealManish** 1 year, 7 months ago

Thanks for helping me look at this again. I tried all of the options

-d : created the /comptia/projects folder and set user folder to /comptia/projects

-m : useradd -m is not valid.. even though -m is an option, it does not seem to work solo

-b : this created a folder called user2 in /comptia/projects and assigned that as the user homedir

-s : sets the shell

B doesn't work. so that brings us back to what is being asked here. is it asking us to assign the user home dir to /comptia/projects? then it's a.
if it's asking to assign the homedir to /comptia/projects/user2, then it's c

upvoted 2 times

  **Veteran903** 1 year, 6 months ago

its C brother, run "man useradd" in your terminal and read on the -b option

upvoted 1 times

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

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

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

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

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

Suggested Answer: B

Community vote distribution

B (100%)

Alizadeh 10 months, 2 weeks ago

Selected Answer: B

The correct answer is B. Replace the failed drive and reconfigure the mirror.

upvoted 1 times

linux_admin 1 year, 4 months ago

Selected Answer: B

To recover an LVM-mirrored volume after a failure in one of its physical volumes, the systems administrator should replace the failed drive and reconfigure the mirror. Therefore, the correct option is B.

LVM mirroring provides redundancy by creating an identical copy of a logical volume on another physical volume. If one of the physical volumes fails, the system can continue to function using the copy on the remaining volume. To recover the volume, the failed physical volume should be replaced and the mirror should be reconfigured to ensure redundancy is restored. Depending on the configuration of the volume, the system may need to be rebooted to recognize the new physical volume.

upvoted 2 times

Ckl22 1 year, 6 months ago

Selected Answer: B

Answer is B

To rebuild the mirrored volume, you replace the broken drive and recreate the physical volume

https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/6/html/logical_volume_manager_administration/mirrorrecover

upvoted 1 times

A systems administrator created a new Docker image called test. After building the image, the administrator forgot to version the release. Which of the following will allow the administrator to assign the v1 version to the image?

- A. docker image save test test:v1
- B. docker image build test:vl
- C. docker image tag test test:vl
- D. docker image version test:v1

Suggested Answer: C

Community vote distribution

C (100%)

🗲️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: C

The correct answer is C. docker image tag test test:v1.

upvoted 1 times

🗲️ 👤 **linux_admin** 1 year, 4 months ago

Selected Answer: C

To assign a version to a Docker image that has already been built and tagged, the systems administrator should use the docker image tag command. Therefore, the correct option is C, docker image tag test test:v1.

The docker image tag command is used to assign a new tag (including a version number) to an existing Docker image.

upvoted 2 times

🗲️ 👤 **Ckl22** 1 year, 6 months ago

Selected Answer: C

Answer is C

\$ docker image tag SOURCE_IMAGE[:TAG] TARGET_IMAGE[:TAG]

https://docs.docker.com/engine/reference/commandline/image_tag/

upvoted 1 times

A Linux systems administrator receives a notification that one of the server's filesystems is full. Which of the following commands would help the administrator to identify this filesystem?

- A. lsblk
- B. fdisk
- C. df -h
- D. du -ah

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: C

The correct answer is C. df -h.

upvoted 2 times

🗳️ 👤 **BreakOff874** 1 year, 3 months ago

Why not D?

upvoted 2 times

🗳️ 👤 **Alizadeh** 10 months, 2 weeks ago

Option D, du -ah, will recursively calculate the disk space usage of all files and directories. This is not the correct command to use to identify a filesystem that is full.

upvoted 2 times

🗳️ 👤 **linux_admin** 1 year, 4 months ago

Selected Answer: C

To identify the filesystem that is full, the Linux systems administrator can use the df -h command. Therefore, the correct option is C, df -h.

The df (disk free) command is used to report the amount of free and used disk space on a filesystem. The -h option displays the output in a human-readable format, making it easier to read.

upvoted 3 times

A systems administrator is notified that the mysqld process stopped unexpectedly. The systems administrator issues the following command:

```
sudo grep -i -r 'out of memory' /var/log
```

The output of the command shows the following:

kernel: Out of memory: Kill process 9112 (mysqld) score 511 or sacrifice child.

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

- A. free -h
- B. nc -v 127.0.0.1 3306
- C. renice -15 \$(pidof mysql)
- D. lsblk
- E. killall -15
- F. vmstat -a 1 4

Suggested Answer: AC

Community vote distribution

AF (100%)

 **TheRealManish**  1 year, 7 months ago

Selected Answer: AF

A: i like A because it will tell us if we need to add more swap space.

B: this is netcat command so not relevant

C: we might need to renice, in the event that a child process was killed, but we dont know if it was the child process or not. but regargles, -15 is not a very nice number.


D: listing storage doesnt help really

E: invalid syntax

F: not sure how this info will help.


By process of elimination I am going with A to check if we can add some swap.. and F to get an overall fee for how the system is doing.. this question is ridiculus

upvoted 5 times

 **Nvoid** 1 year, 7 months ago

Good effort, keep it up buddy!

upvoted 1 times

 **Nvoid** 1 year, 7 months ago

no idea if we're really correct, but just gonna keep trying.

upvoted 1 times

 **dsmitd33**  6 months, 2 weeks ago

Selected Answer: AF

is vmstat -a 1 4 even a valid command?

upvoted 1 times

 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: AF

The correct answers are A. free -h and F. vmstat -a 1 4

upvoted 1 times

 **angellorv** 1 year, 2 months ago

Answer C is correct:

as the administrator you are trying to prevent a database crash caused by OOM Killer. Questions with OOM caused by MySQL consider "renice" - reprioritizing MySQL as the leading answer.

upvoted 1 times

🗨️ 👤 **linux_admin** 1 year, 4 months ago

Selected Answer: AF

To troubleshoot an out-of-memory issue with the mysql process, the systems administrator should check the server's memory usage. Therefore, the correct option is A, free -h, and F, vmstat -a 1 4.

upvoted 1 times

🗨️ 👤 **linux_admin** 1 year, 4 months ago

Option B (nc -v 127.0.0.1 3306) is used to check if the MySQL server is running and accepting connections. It will not help troubleshoot an out-of-memory issue.

Option C (renice -15 \$(pidof mysql)) is used to change the priority of a running process. It will not help troubleshoot an out-of-memory issue.

Option D (lsblk) is used to list information about all available block devices, including their mount points. It will not help troubleshoot an out-of-memory issue.

Option E (killall -15) is used to send a signal to all processes with a particular name. It will not help troubleshoot an out-of-memory issue.

upvoted 2 times

🗨️ 👤 **KnifeClown1** 1 year, 4 months ago

Selected Answer: AF

A. free -h

F. vmstat -a 1 4

upvoted 2 times

🗨️ 👤 **Nvoid** 1 year, 7 months ago

Selected Answer: AF

We can eliminate B/C/D because they have nothing to do with it, which leaves us to `troubleshoot with A/E/F`

We eliminate E because that is killing processes, not troubleshooting.

We're left with A & F.

A: Troubleshoot memory = yes

B: open netcat on port 3306 = no

C: renice -15 = no

D: show block devices = no

E: kill all processes in -15 = no

F: Troubleshoot memory = yes

upvoted 2 times

🗨️ 👤 **Veteran903** 1 year, 7 months ago

Well, analyzing this one again.... I'll go with EF just because the other answers dont match and like you guys said an stopped process cant be renice, also option A offers just info about system memory and swap memory so i think the vmstat is a way better option...

upvoted 1 times

🗨️ 👤 **Veteran903** 1 year, 7 months ago

Hello guys, the right answers are C and F, A free -h only provides info about system memory and swap space while vmstat reports virtual memory statistics, the reason why you have to pick C is because you are not killing the process itself but "sacrificing the child", doing this you will about killing or stopping the process, CompTIA wording is brutal!!!! lol

upvoted 2 times

🗨️ 👤 **MissAllen** 1 year, 7 months ago

I would say AC are the correct answers. We don't know the current nice value of the mysql process, plus since it stopped how can you renice it?

Renice is for running processes. Renice also affect CPU access, not memory. Vmstat would be a better choice.

upvoted 1 times

🗨️ 👤 **TheRealManish** 1 year, 7 months ago

wait, you say AC is correct, but you can't renice a stopped process. how is C correct then? thanks

upvoted 1 times

🗨️ 👤 **Alizadeh** 10 months, 2 weeks ago

You are correct, you cannot renice a stopped process. The renice command only works on running processes.

In this case, the mysql process is stopped, so the renice -15 \$(pidof mysql) command will not have any effect. However, the administrator may still want to run the command to see if it helps to resolve the issue.

upvoted 1 times

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

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

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ **linux_admin** 10 months, 2 weeks ago

Selected Answer: B

o restore a Linux server to its usual target after being switched to rescue.target mode, the administrator can use the following command:

C. systemctl get-default

This command will display the current default target for the system. If the default target was changed to rescue.target for maintenance, the output of the command will be rescue.target.

upvoted 1 times

🗳️ **linux_admin** 10 months, 2 weeks ago

Selected answer is C

upvoted 1 times

🗳️ **linux_admin** 10 months, 1 week ago

Picking B.

When the system is restarted, the systemd init system will automatically start the services required for the default target. This will include services such as the graphical user interface, networking, and other essential services.

upvoted 1 times

🗳️ **Nvoid** 1 year, 1 month ago

Selected Answer: B

A: PowerDown (state 0)

B: reboot server

C: show default target

D: not sure this is even a command.

Picking C here.

upvoted 1 times

🗳️ **Nvoid** 1 year, 1 month ago

Opps, meant say "Picking B here." sawery ppls.

upvoted 2 times

A systems administrator was tasked with assigning the temporary IP address/netmask 192.168.168.1/255.255.255.255 to the interface eth0 of a Linux server.

When adding the address, the following error appears:

```
# ip address add 192.168.168.1/33 dev eth0
```

Error: any valid prefix is expected rather than "192.168.168.1/33".

Based on the command and its output above, which of the following is the cause of the issue?

- A. The CIDR value /33 should be /32 instead.
- B. There is no route to 192.168.168.1/33.
- C. The interface eth0 does not exist.
- D. The IP address 192.168.168.1 is already in use.

Suggested Answer: C

Community vote distribution

A (100%)

  **ryanzou** Highly Voted 2 years, 1 month ago

Selected Answer: A

Never see there is a subnet mask of /33



upvoted 7 times

  **Nvoid** Highly Voted 2 years, 1 month ago

Selected Answer: A



Mod Fix. Thx!!!

upvoted 5 times

  **Nvoid** 2 years, 1 month ago

Mod Fix. it should be A.

upvoted 1 times

  **Nvoid** 2 years, 1 month ago

Mod still needs to fix.


upvoted 1 times

  **tegami** Most Recent 11 months, 2 weeks ago

Selected Answer: A

Oh common guys it's A

upvoted 1 times

  **Alizadeh** 1 year, 4 months ago

Selected Answer: A

The correct answer is A. The CIDR value /33 should be /32 instead.

upvoted 1 times

  **linux_admin** 1 year, 10 months ago

Selected Answer: A

The cause of the issue is that the prefix length specified in the IP address/netmask is invalid. In the command `ip address add 192.168.168.1/33 dev eth0`, the prefix length is specified as /33, which is not a valid prefix length.

In IP networking, the prefix length specifies the number of bits in the network portion of the address. For example, a netmask of 255.255.255.0 corresponds to a prefix length of /24, because the first 24 bits of the address are used for the network portion.


The valid prefix lengths depend on the address class and type. For IPv4 addresses, the prefix length must be between 0 and 32. A prefix length of 33 is not valid, and that is why the command produced an error.

To fix the issue, the systems administrator should use a valid prefix length in the IP address/netmask, such as /32 if the address is a host address or

/24 if the address is a network address. For example:

```
ip address add 192.168.168.1/32 dev eth0
```

upvoted 2 times

  **Pinnubhai** 1 year, 11 months ago

Selected Answer: A

AAAAAAA

upvoted 5 times

A Linux user reported the following error after trying to connect to the system remotely: ssh: connect to host 10.0.1.10 port 22: Resource temporarily unavailable

The Linux systems administrator executed the following commands in the Linux system while trying to diagnose this issue:

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

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

Which of the following commands will resolve this issue?

- A. firewall-cmd --zone=public --permanent --add-service=22
- B. systemctl enable firewalld; systemctl restart firewalld
- C. firewall-cmd --zone=public --permanent --add-service=ssh
- D. firewall-cmd --zone=public --permanent --add-port=22/udp

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Qubert2** 8 months, 1 week ago

C but it won't work until the firewall is reloaded with: firewall-cmd --reload
upvoted 1 times

🗳️ 👤 **linux_admin** 2 years, 4 months ago

Selected Answer: C

The command firewall-cmd --zone=public --permanent --add-service=ssh is used to add the SSH service to the list of allowed services in the firewall configuration of a Linux system.

Here is a breakdown of the command:

firewall-cmd is the command-line tool used to manage the firewall configuration in a Linux system that uses the firewalld firewall daemon.

--zone=public specifies the firewall zone to which the rule should be applied. The public zone is one of the predefined zones in firewalld that is typically used for public-facing network interfaces.

--permanent specifies that the rule should be made persistent across firewall reloads or system reboots.



--add-service=ssh specifies that the ssh service should be allowed in the firewall configuration. The ssh service is a predefined service in firewalld that allows incoming SSH traffic on port 22.

upvoted 3 times

🗳️ 👤 **linux_admin** 2 years, 4 months ago

When the command is executed, the firewall configuration is updated to allow incoming SSH traffic on port 22 for the public zone, and the configuration change is made persistent across firewall reloads or system reboots. This allows SSH connections to be established with the system from remote hosts, provided that the SSH service is running on the system and the SSH port is not blocked by any other firewall or security measure.

upvoted 3 times

  **Nvoid** 2 years, 7 months ago

Selected Answer: C

https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/security_guide/sec-working_with_zones

upvoted 2 times

A Linux administrator has been tasked with installing the most recent versions of packages on a RPM-based OS. Which of the following commands will accomplish this task?

- A. apt-get upgrade
- B. rpm -a
- C. yum updateinfo
- D. dnf update
- E. yum check-update

Suggested Answer: D

Community vote distribution

D (100%)

🗨️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: D

D. dnf update.

upvoted 2 times

🗨️ 👤 **linux_admin** 1 year, 4 months ago

Selected Answer: D

D. dnf update

Explanation:

dnf is the default package manager for RPM-based Linux distributions like Fedora, CentOS, and RHEL. The update command is used to update all installed packages to their latest versions. Running dnf update will download and install the latest versions of all installed packages, along with any necessary dependencies.

The other options are not correct for the following reasons:

apt-get is a package manager for Debian-based distributions and not RPM-based distributions.

rpm -a is used to query information about all installed packages and not update them.

yum updateinfo is used to display information about available updates but not update packages.

yum check-update is used to check for available updates but not install them.

upvoted 2 times

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

- A. partprobe
vgcreate
lvextend
- B. lvcreate
fdisk
partprobe
- C. fdisk
partprobe
mkfs
- D. fdisk
pvcreate
vgextend

Suggested Answer: D

Community vote distribution

D (100%)

🗨️ **Alizadeh** 10 months, 2 weeks ago

Selected Answer: D

The correct sequence of commands is D. fdisk, pvcreate, vgextend.
upvoted 2 times

🗨️ **linux_admin** 1 year, 4 months ago

Selected Answer: D

D. fdisk, partprobe, pvcreate, vgextend

Explanation:

fdisk is used to create a partition on the new disk.

partprobe is used to inform the OS of the new partition and update the kernel's partition table.

pvcreate is used to create a physical volume on the new partition.

vgextend is used to extend the volume group to include the new physical volume.

The other options are not correct for the following reasons:

Option A is not correct because it is missing the pvcreate and vgextend commands, which are required to create a physical volume and extend the volume group to include the new physical volume.

Option B is not correct because it uses lvcreate instead of pvcreate to create a new logical volume, which is not necessary for expanding a volume group using a new disk.

Option C is not correct because it uses mkfs to create a file system, which is not necessary for expanding a volume group using a new disk.

Additionally, it is missing the pvcreate and vgextend commands, which are required to create a physical volume and extend the volume group to include the new physical volume.

upvoted 3 times

Which of the following directories is the mount point in a UEFI system?

- A. /sys/efi
- B. /boot/efi
- C. /efi
- D. /etc/efi

Suggested Answer: B

Community vote distribution

B (100%)

🗉 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: B

The correct answer is B. /boot/efi.

upvoted 1 times

🗉 👤 **linux_admin** 1 year, 4 months ago

Selected Answer: B

In a UEFI (Unified Extensible Firmware Interface) system, the system firmware requires a partition to store boot loaders and other boot-related data. This partition is usually a small FAT32 partition, and it is mounted as the EFI system partition (ESP). The most common mount point for the ESP is /boot/efi. It is usually a small partition with a size between 100MB and 500MB and it is shared between different operating systems in a multi-boot environment. The contents of the ESP typically include boot loaders, kernel images, configuration files, and other boot-related files

upvoted 3 times

🗉 👤 **abrilo** 1 year, 6 months ago

Selected Answer: B

the ESP setup utilizes the old microsoft file allocation table (FAT) filesystem to store the bootloader programs. on linux systems, the ESP is typically mounted in the /boot/efi directory....

upvoted 1 times

A Linux administrator copied a Git repository locally, created a feature branch, and committed some changes to the feature branch. Which of the following Git actions should the Linux administrator use to publish the changes to the main branch of the remote repository?

- A. rebase
- B. tag
- C. commit
- D. push

Suggested Answer: C

Community vote distribution

D (100%)

🗳️ 👤 **Qubert2** 8 months, 1 week ago

Selected Answer: D

There are several steps to publish local, committed changes on a local branch to the remote main branch. The final command is "push", so D is correct, but it's a bad question.

upvoted 1 times

🗳️ 👤 **HMAC** 2 years, 1 month ago

Since "committed some changes to the feature branch" was present in the question. Push is the next step.

upvoted 1 times

🗳️ 👤 **linux_admin** 2 years, 4 months ago

Selected Answer: D

In Git, the "push" command is used to publish local changes to a remote repository. When the administrator makes changes to the feature branch, the changes are only available in the local repository. If the changes are ready to be merged into the main branch of the remote repository, the administrator can use the "push" command to upload the changes to the remote repository. This will make the changes available to others who are also working on the same repository.

Before pushing the changes to the remote repository, the administrator should first ensure that the changes are committed to the local repository using the "commit" command. The administrator should also ensure that the local repository is up-to-date with the remote repository by using the "pull" command to download any changes made by others before attempting to push the changes. Once the changes are committed and the local repository is up-to-date, the administrator can use the "push" command to upload the changes to the remote repository.

upvoted 3 times

🗳️ 👤 **KnifeClown1** 2 years, 4 months ago

Selected Answer: D

D. push

upvoted 1 times

🗳️ 👤 **ethan_me** 2 years, 7 months ago

Selected Answer: D

Commit it for local changes. Push is for remote changes.

upvoted 4 times

🗳️ 👤 **Veteran903** 2 years, 7 months ago

The amount of wrong answers is staggering, Im stating to doubt if these are real examn questions

upvoted 3 times

🗳️ 👤 **Nvoid** 2 years, 7 months ago

in my experience passing 4 other exams using ET, if the questions are bad or have errors from what the original should/could be, it probably won't be on the test. The solid ones that are no question usually on legit.

upvoted 3 times

🗳️ 👤 **Veteran903** 2 years, 6 months ago

Good to know, make sense, thanks

upvoted 1 times

A Linux administrator needs to obtain a list of all volumes that are part of a volume group. Which of the following commands should the administrator use to accomplish this task?

- A. vgs
- B. lvs
- C. fdisk -l
- D. pvs

Suggested Answer: B

Community vote distribution

B (71%)

A (29%)

🗳️ 👤 **johnsie** 1 month ago

Selected Answer: B

Drvision below nailed it
upvoted 1 times

🗳️ 👤 **bongobo** 9 months, 2 weeks ago

L ist V olume S
upvoted 4 times

🗳️ 👤 **DRVision** 1 year ago

Selected Answer: B

The lvs command in Linux is used to display information about logical volumes. Here's what the other options do:

A. vgs: This command provides information about volume groups. C. fdisk -l: This is not a valid command. Perhaps you meant fdisk -l, which lists all partitions on all hard drives. D. pvs: This command provides information about physical volumes.

So, for your specific task, lvs is the appropriate command.
upvoted 1 times

🗳️ 👤 **DRVision** 1 year ago

The command vgs is used to display information about volume groups, not the volumes that are part of a volume group. It provides a summary of available and used space in a volume group but does not list out the individual logical volumes within that group.

On the other hand, the lvs command lists all logical volumes in all volume groups if no arguments are given, or all logical volumes in the specified volume group. This makes lvs the correct command to use when you want to see a list of all volumes that are part of a specific volume group.

So, while vgs is a useful command for getting an overview of volume groups, it doesn't provide the level of detail (i.e., the individual volumes) that the question is asking for. That's why the answer is not A.
upvoted 2 times

🗳️ 👤 **tutita** 1 year, 7 months ago

Selected Answer: B

answer b lvs, the question states "logical volumes that are part of a group volume"
upvoted 2 times

🗳️ 👤 **linux_admin** 1 year, 10 months ago

Selected Answer: B



The Linux administrator should use the lvs command to obtain a list of all volumes that are part of a volume group. The lvs command is used to display information about logical volumes in the LVM. Running the command without any options will provide a summary of all the logical volumes, including the name of the volume group and the size of each volume. By providing the name of the volume group as an argument to the lvs command, the administrator can obtain a list of all volumes that are part of the volume group.
upvoted 1 times

🗳️ 👤 **Pinnubhai** 1 year, 11 months ago

Selected Answer: B

B is right because it shows a list of all logical volumes and a column VG of which vg they are part of.

upvoted 1 times

  **Nvoid** 2 years, 1 month ago

Selected Answer: A



Answer is A, VGS is for volume group. LVS is or logical Volumes.

upvoted 2 times

  **TheRealManish** 2 years, 1 month ago



The wording on this one has confused me. "obtain a list of all volumes that are part of a volume group. Well, as we know physical volumes make up volume groups which then are carved into logical volumes. so the questions is, when it asks "obtain a list of all volumes part of a volume group" does it mean a list of all physical volumes? Does it mean a list of all physical volumes? or all logical volumes? VGS is going to show you a list of volume groups, but won't show you which logical of physical volumes are involved. so I argue that It's not A... but instead either LVS to list the logical volumes that are part of the volume groups.. or PVS to list the physical volumes that are part of the volume group. the wording is terrible

upvoted 1 times

  **Nvoid** 2 years, 1 month ago



damn, thats interesting, i may have to change my answer last minute. i figured VGS would show all the logical volumes, but ya it probably won't.

upvoted 1 times

  **Nvoid** 2 years, 1 month ago

ya, i'm not sure what to select now, but its got to be one or the other .. 50/50. lol

upvoted 1 times

  **Nvoid** 2 years, 1 month ago

Mod Fix, it's A.

upvoted 1 times

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

- A. pull -> push -> add -> checkout
- B. pull -> add -> commit -> push
- C. checkout -> push -> add -> pull
- D. pull -> add -> push -> commit

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **johnsie** 1 month ago

Selected Answer: B

i concur with nixonbii
upvoted 1 times

🗳️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: B

The correct answer is B. pull -> add -> commit -> push.
upvoted 2 times

🗳️ 👤 **nixonbii** 1 year, 4 months ago

This question really needs to clarify if the file is being added to a REMOTE or LOCAL repository. It matters.
upvoted 2 times

🗳️ 👤 **linux_admin** 1 year, 4 months ago

Selected Answer: B

The correct order of Git commands to add a new configuration file to a Git repository is:

pull to get the latest changes from the remote repository.

add to add the new configuration file to the staging area.

commit to commit the changes to the local repository.

push to push the changes to the remote repository.

Therefore, the correct answer is:

B. pull -> add -> commit -> push

upvoted 2 times

A systems administrator is tasked with mounting a USB drive on a system. The USB drive has a single partition, and it has been mapped by the system to the device `/dev/sdb`. Which of the following commands will mount the USB to `/media/usb`?

- A. `mount /dev/sdb1 /media/usb`
- B. `mount /dev/sdb0 /media/usb`
- C. `mount /dev/sdb /media/usb`
- D. `mount -t usb /dev/sdb1 /media/usb`

Suggested Answer: A

Community vote distribution

A (100%)

🗨️ 👤 **Alizadeh** 10 months, 2 weeks ago

Selected Answer: A

The correct answer is A. `mount /dev/sdb1 /media/usb`.

upvoted 2 times

🗨️ 👤 **linux_admin** 1 year, 4 months ago

Selected Answer: A

A. `mount /dev/sdb1 /media/usb`

Explanation:

The "mount" command is used to mount file systems, and in this case, we are using it to mount the USB drive. The device name for the USB drive is `/dev/sdb`, and since it has only one partition, the partition number is 1. Therefore, we need to mount `/dev/sdb1` to the desired mount point, which is `/media/usb`. The correct command for this would be "mount `/dev/sdb1 /media/usb`". Option B is incorrect because the partition number is 1, not 0. Option C is incorrect because we need to specify the partition (`/dev/sdb1`) and not just the device (`/dev/sdb`). Option D is incorrect because "-t usb" is not a valid option for the mount command.

upvoted 3 times

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

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

Suggested Answer: C

Community vote distribution

A (100%)

🗳️ 👤 **linux_admin** Highly Voted 1 year, 4 months ago

Selected Answer: A

The correct option is A: `chown user2:accounting script.sh` `chmod 750 script.sh`.

Explanation: The `chown` command will change the owner of the file to `user2` and the group to `accounting`. The `chmod` command will change the file permissions to `750`, which will allow the owner (`user2`) to read, write, and execute the file, members of the `accounting` group to read and execute the file, and deny all other users from accessing the file. This configuration will allow members of the `accounting` group to execute the script but not make any changes to it, which is the desired outcome.

upvoted 5 times

🗳️ 👤 **Rob74613** Most Recent 1 year ago

Selected Answer: A

For sure A

Keeps full control to User2

Grants r-x (read and execute) to the accounting group

Denies any permissions to all other users

upvoted 2 times

🗳️ 👤 **KnifeClown1** 1 year, 4 months ago

Selected Answer: A

A, no doubt

upvoted 1 times

🗳️ 👤 **wajeed_sulu** 1 year, 5 months ago

Selected Answer: A

I believe A is the correct answer with proper access permissions.

upvoted 1 times

🗳️ 👤 **MissAllen** 1 year, 7 months ago

Answer A is correct. `Chown` command syntax is:

`chown username:groupname filename`

Adding the accounting group to the file and granting `rxr-x---` to the file provides the proper access.

upvoted 4 times

🗳️ 👤 **Veteran903** 1 year, 7 months ago

exactly! thanks

upvoted 2 times

🗳️ 👤 **Nvoid** 1 year, 7 months ago

A, way to go ppl!

upvoted 1 times

A systems administrator needs to verify whether the built container has the app.go file in its root directory. Which of the following can the administrator use to verify the root directory has this file?

- A. docker image inspect
- B. docker container inspect
- C. docker exec <container_name> ls
- D. docker ps <container_name>

Suggested Answer: B

Community vote distribution

C (100%)

linux_admin **Highly Voted** 1 year, 4 months ago

Selected Answer: C

C. docker exec <container_name> ls

This command runs the ls command inside the specified container and lists the contents of the current directory. The administrator can use this to verify whether the app.go file is present in the container's root directory.

Option B (docker container inspect) would provide detailed information about the specified container, such as its configuration and state, but it would not show the contents of the container's file system. To view the files inside the container, the administrator should use the docker exec command to execute a shell command inside the container, such as ls to list the files in the root directory.

upvoted 5 times

Alizadeh **Most Recent** 10 months, 2 weeks ago

Selected Answer: C

The correct answer is C. docker exec <container_name> ls.

upvoted 4 times

Amazing475 11 months, 3 weeks ago

C is the correct option

upvoted 1 times

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

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

The Linux administrator is trying to select the appropriate syntax formatter to correct any issues with the configuration file. Which of the following should the syntax formatter support to meet this goal?

- A. Markdown
- B. XML
- C. YAML. JSON

Suggested Answer: C

Community vote distribution

C (100%)

linux_admin 10 months, 2 weeks ago

Selected Answer: C

C. YAML

upvoted 4 times

Lwarder1 10 months, 3 weeks ago

Mistyped Answers:

- A. Markdown
- B. XML
- C. YAML
- D. JSON

Answer: C. YAML

upvoted 3 times

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

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

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

- A. Comment out the /opt/app filesystem in /etc/fstab and reboot.
- B. Reformat the /opt/app filesystem and reboot.
- C. Perform filesystem checks on local filesystems and reboot.
- D. Trigger a filesystem relabel and reboot.

Suggested Answer: A

Community vote distribution

A (100%)

 **linux_admin** Highly Voted 10 months, 2 weeks ago

Selected Answer: A

Option A is the most appropriate choice because it will prevent the /opt/app filesystem from being mounted during boot, allowing the system to boot normally. The error message shows that the mount point for /opt/app is failing, and commenting out the entry in the /etc/fstab file will prevent the system from trying to mount the filesystem during boot. This is a quick solution to allow the system to boot normally and does not require any additional troubleshooting or repair.

upvoted 5 times

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

Output 1:

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

Output 2:

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

Output 3:

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

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

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

Suggested Answer: A

Community vote distribution

D (75%)

C (25%)

MissAllen Highly Voted 2 years, 7 months ago

I choose answer D. The system killed the java process because of lack of memory, but no one has configured a swap partition! That will help. A swap partition is generally recommended as one of the two main partitions created during a linux installation.

upvoted 11 times

Veteran903 2 years, 7 months ago

Im with you, its D

upvoted 3 times

IFBBPROSALCEDO Most Recent 11 months, 2 weeks ago

Selected Answer: C

C. Configure more CPU cores to allow for the server to allocate more processing and prevent the Java process from consuming all of the available resources.

Adding more CPU cores can help distribute the load and handle higher CPU demands, potentially resolving the issue of high CPU usage by the Java process.

Adding more CPU cores can help distribute the load and handle higher CPU demands, potentially resolving the issue of high CPU usage by the Java process.

upvoted 1 times



DRVision 1 year, 6 months ago

Selected Answer: D

Given this, the best solution would be to configure the swap space to allow for spikes in usage during peak hours and prevent the Java process from stopping due to a lack of memory. This can be done by increasing the size of the swap space or adding more swap space. This will provide a buffer for the system when it runs out of physical memory, preventing the OOM killer from being invoked and the Java process from being killed. So, the

correct answer is D. Configure the swap space to allow for spikes in usage during peak hours and prevent the Java process from stopping due to a lack of memory.

upvoted 3 times

  **BreakOff874** 2 years, 3 months ago

Selected Answer: D

the last output shows that there is no swap available


upvoted 1 times

A Linux administrator found many containers in an exited state. Which of the following commands will allow the administrator to clean up the containers in an exited state?

- A. `docker rm --all`
- B. `docker rm $(docker ps -aq)`
- C. `docker images prune *`
- D. `docker rm --state exited`

Suggested Answer: D

Community vote distribution



🗳️ 👤 **Mistermiyagi** 2 months, 3 weeks ago

Selected Answer: B

For those confused on B or C. The question states containers not images. C would be for images and B would be for containers.
upvoted 1 times

🗳️ 👤 **Qubert2** 8 months, 1 week ago

In the real world, you would use: `docker container prune`
upvoted 2 times

🗳️ 👤 **DRVision** 1 year, 6 months ago

Selected Answer: B

I think B has a typo, it should be `$(docker ps -aq --filter status=exited)` otherwise you removed all the containers
upvoted 4 times

🗳️ 👤 **salthedhash** 1 year, 7 months ago

Selected Answer: B

B is the answer
upvoted 1 times

🗳️ 👤 **Amazing475** 1 year, 11 months ago

Selected Answer: B

B is the correct answer
upvoted 2 times

🗳️ 👤 **LKRISB** 2 years, 1 month ago

B. `docker rm $(docker ps -aq)`

Explanation:

The command `docker ps` is used to list running containers.

The option `-a` is used to list all containers, including those in an exited state.

The option `-q` is used to display only the container IDs.

The command `docker rm` is used to remove containers.

By using `$(docker ps -aq)` as a parameter to `docker rm`, it will remove all containers returned by the `docker ps -aq` command, which includes containers in an exited state.

upvoted 4 times

🗳️ 👤 **POGActual** 2 years, 3 months ago

It unfortunately cant be B, as that would remove all container.

"You can review the containers on your system with `docker ps`. Adding the `-a` flag will show all containers. When you're sure you want to delete them, you can add the `-q` flag to supply the IDs to the `docker stop` and `docker rm` commands:

`docker stop $(docker ps -a -q)`

`docker rm $(docker ps -a -q)`

<https://www.digitalocean.com/community/tutorials/how-to-remove-docker-images-containers-and-volumes>

upvoted 2 times

🗨️ **linux_admin** 2 years, 4 months ago

Selected Answer: D

The difference between `docker rm $(docker ps -aq)` and `docker rm --state exited` is that the first command will remove all containers, regardless of their state, while the second command will only remove containers in an exited state.

`docker rm $(docker ps -aq)` will remove all containers because the `docker ps -aq` command lists all containers' IDs in quiet mode, and then `docker rm` removes them. This command will remove all containers, including those that are running, stopped, and exited.

`docker rm --state exited`, on the other hand, only removes containers in an exited state. The `--state` option is used to specify the state of the containers to be removed, and `exited` is the state that we want to remove. This command is useful for cleaning up containers that have exited and are no longer needed, freeing up resources and disk space occupied by these containers.

In summary, `docker rm $(docker ps -aq)` removes all containers, while `docker rm --state exited` removes only containers in an exited state.

upvoted 1 times

🗨️ **linux_admin** 2 years, 4 months ago

Of the options given, only option D, `docker rm --state exited`, would work, but it has a syntax error, as you pointed out. The correct syntax for removing containers in an exited state using Docker is:

```
docker rm $(docker ps -aq --filter "status=exited")
```

This command filters the list of containers to those that are in an exited state, and then removes them.

upvoted 1 times

🗨️ **KnifeClown1** 2 years, 4 months ago

Selected Answer: B

B. `docker rm $(docker ps -aq)`

upvoted 1 times

🗨️ **Ckl22** 2 years, 6 months ago

Selected Answer: B

B is the answer

"The `docker ps` command only shows running containers by default. To see all containers, use the `-a` (or `--all`) flag"

"The example below uses `docker ps -q` to print the IDs of all containers that have exited (`--filter status=exited`), and removes those containers with the `docker rm` command"

<https://docs.docker.com/engine/reference/commandline/docker/>

upvoted 1 times

🗨️ **Nvoid** 2 years, 7 months ago

Selected Answer: B

going with B, Ref: https://coderwall.com/p/zguz_w/docker-remove-all-exited-containers

upvoted 1 times

🗨️ **MaryamNesa** 2 years, 7 months ago

Correct answer is B

upvoted 2 times

🗨️ **TheRealManish** 2 years, 7 months ago

I'm googling and i see that : B: `docker rm $(docker ps -aq)` will list and remove all docker containers.. `-a` is list `-q` to pass it to the `rm` command. you need the `-f status=exited` in that command for it to be valid. As for D, that command doesnt even work at all. so i dont konw what to think. B seems the closest though it's also not right :(

upvoted 1 times

🗨️ **Veteran903** 2 years, 7 months ago

We better go with B on this one, its the best answer.....

upvoted 2 times

A Linux administrator reviews a set of log output files and needs to identify files that contain any occurrence of the word denied. All log files containing entries in uppercase or lowercase letters should be included in the list. Which of the following commands should the administrator use to accomplish this task?

- A. `find . -type f -print | xargs grep -ln denied`
- B. `find . -type f -print | xargs grep -nv denied`
- C. `find . -type f -print | xargs grep -wL denied`
- D. `find . -type f -print | xargs grep -li denied`

Suggested Answer: A

Community vote distribution

D (100%)

🗳️ 👤 **Amazing475** 11 months, 3 weeks ago

Selected Answer: D

D. `find . -type f -print | xargs grep -li denied`
upvoted 2 times

🗳️ 👤 **linux_admin** 1 year, 4 months ago

Selected Answer: D

D. `find . -type f -print | xargs grep -li denied`

Explanation:

`find . -type f -print` finds all files (-type f) in the current directory and its subdirectories (.) and prints their names (-print).

`xargs` passes the list of file names to the `grep` command as arguments.

`grep -li denied` searches for the word "denied" in the files, ignoring case (-i) and printing only the names of the files that contain the word (-l).

Option A (`grep -ln denied`) would print the name of the file and the line number of the occurrence of "denied", but it doesn't meet the requirement of including both uppercase and lowercase occurrences.

Option B (`grep -nv denied`) would print the name of the file and the line number of all the lines that do not contain "denied", which is not what is required.

Option C (`grep -wL denied`) would print the name of the files that do not contain the exact word "denied" (i.e., it would exclude files that contain "permission denied" or "access denied"), which is also not what is required.

upvoted 3 times

🗳️ 👤 **KnifeClown1** 1 year, 4 months ago

Selected Answer: D

D. `find . -type f -print | xargs grep -li denied`

The command `find . -type f -print | xargs grep -li denied` will accomplish the task. The `find` command is used to search for files in the current directory and all its subdirectories. The `-type f` option limits the search to regular files, and the `-print` option causes the names of the files to be printed. The output from the `find` command is then passed to `xargs`, which is used to execute the `grep` command on each file. The `grep` command is used to search for occurrences of the word "denied". The `-l` option causes `grep` to print only the names of the files that contain a match, and the `-i` option causes `grep` to perform a case-insensitive search, so that both uppercase and lowercase occurrences of the word "denied" will be matched.

upvoted 1 times

🗳️ 👤 **Nvoid** 1 year, 7 months ago

Selected Answer: D

need the -i for ignore case going with `D` !!
upvoted 4 times

🗳️ 👤 **Nvoid** 1 year, 7 months ago

btw its xargs not xrgs, mod please fix.

upvoted 1 times

🗨️ 👤 **SaadiaS** 1 year, 7 months ago

<https://www.geeksforgeeks.org/grep-command-in-unixlinux/>

Voting for D

upvoted 1 times

🗨️ 👤 **MaryamNesa** 1 year, 7 months ago

D is correct

upvoted 4 times

🗨️ 👤 **TheRealManish** 1 year, 7 months ago

@MissAllen , i'm up voting you, because all of my research agrees that -i is to ignore case. However, on my centos VM, when i don't use -i it's still ignoring case by default, weird.

upvoted 2 times

🗨️ 👤 **MissAllen** 1 year, 7 months ago

I disagree. Answer D is better. We need grep -i (ignore case) to ensure we search for upper or lowercase denied.

upvoted 3 times

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

- A. docker builder -f --name web-01 httpd
- B. docker load --name web-01 httpd
- C. docker ps -a --name web-01 httpd
- D. docker run -d --name web-01 httpd

Suggested Answer: D

🗒️ 👤 **linux_admin** Highly Voted 10 months, 2 weeks ago

D. docker run -d --name web-01 httpd

Explanation:

docker run command is used to create and run a container based on a specified image.

The -d option is used to start the container in the background (detached mode).

The --name web-01 option is used to set the name of the container to web-01.

httpd is the name of the image that the container will be based on.

upvoted 5 times

🗒️ 👤 **Eikan** Most Recent 6 months, 1 week ago

Selected Answer: D

-d flag: for running the commands in the background.

upvoted 1 times

A Linux administrator is providing a new Nginx image from the registry to local cache. Which of the following commands would allow this to happen?

- A. docker pull nginx
- B. docker attach nginx
- C. docker commit nginx
- D. docker import nginx

Suggested Answer: A

  **linux_admin** Highly Voted 10 months, 2 weeks ago

A. docker pull nginx

Explanation:

docker pull command is used to pull an image from a registry and store it in the local Docker image cache.

nginx is the name of the image that will be pulled.

upvoted 5 times

  **Eikan** Most Recent 6 months, 1 week ago

Selected Answer: A

docker pull - This command allows you to pull any image which is present in the official registry of docker, Docker hub. By default, it pulls the latest image, but you can also mention the version of the image.

upvoted 1 times

Which of the following tools is BEST suited to orchestrate a large number of containers across many different servers?

- A. Kubernetes
- B. Ansible
- C. Podman
- D. Terraform

Suggested Answer: A

Community vote distribution

A (100%)

🗉 👤 **linux_admin** 10 months, 2 weeks ago

Selected Answer: A



Kubernetes is a container orchestration platform that automates the deployment, scaling, and management of containerized applications. It is designed to work with large numbers of containers and servers, providing features such as load balancing, automatic scaling, self-healing, and rolling updates. Kubernetes provides a flexible and scalable platform for managing containerized workloads, whether they are running on-premises or in the cloud.

upvoted 3 times

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

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

Suggested Answer: C

  **post20** 9 months, 2 weeks ago

Correct answer option C. This command queries the firewall daemon (firewalld) to check if the "http" service is currently enabled in the firewall. If the service is enabled, the command will return "yes"; otherwise, it will return "no".

upvoted 3 times

A systems administrator is encountering performance issues. The administrator runs a command with the following output:

09:10:18 up 457 days, 32min, 5 users, load average: 4.22 6.63 5.58

The Linux server has the following system properties:

CPU 4 vCPU -

Memory: 50GB -

Which of the following accurately describes this situation?

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

Suggested Answer: B

Community vote distribution

A (63%)

B (38%)

☐ **MissAllen** Highly Voted 2 years, 7 months ago

Disagree. The system **REQUIRES** a reboot? I have seen plenty of servers running for more than a year. The uptime command shows that even with 4 vCPUs they are overloaded. Answer A is better.

upvoted 10 times

☐ **SaadiaS** 2 years, 7 months ago

<https://www.digitalocean.com/community/tutorials/load-average-in-linux>

upvoted 5 times

☐ **SaadiaS** 2 years, 7 months ago

I will go for A as well

upvoted 5 times

☐ **HappyDay030303** Most Recent 2 months, 3 weeks ago

Selected Answer: A

A. values shown (4.22, 6.63, 5.58) exceed the total available CPU capacity (4)

upvoted 1 times

☐ **044f354** 9 months, 2 weeks ago

Selected Answer: A

easy answer.

already explained well by many of you

upvoted 1 times

☐ **DRVision** 1 year, 6 months ago

Selected Answer: A

The output indicates that the system is under CPU pressure. The load average numbers represent the system load time over a period of 1, 5, and 15 minutes. In this case, all three numbers are higher than the number of vCPUs available (4), which suggests that the CPU is a bottleneck. Therefore, the correct answer is A. The system is under CPU pressure and will require additional vCPUs.

Here's a brief explanation of the other options:

B. The system has been running for over a year and requires a reboot. While it's true that the system has been running for over a year, this doesn't necessarily mean it requires a reboot. Linux systems are capable of running for many years without needing a reboot.

C. Too many users are currently logged in to the system. The number of users currently logged in (5) should not be a problem unless they are all running processes that are heavily consuming resources.

D. The system requires more memory. There's no information given that suggests the system is low on memory. The performance issue seems to be related to the CPU, not memory.

upvoted 2 times

🗨️ 👤 **MiraGod** 1 year, 9 months ago

Selected Answer: A

The devices has more load on CPU than CPU Cores, only with that information we know there is an issue.

With the system being up more than a year sure it does not sound nice but only with that information there is no evidence that is the problem

upvoted 1 times

🗨️ 👤 **tutita** 2 years, 1 month ago

Selected Answer: A

option A is the right one, the system is overloaded, doing a reboot wont solve it.

upvoted 1 times

🗨️ 👤 **POGActual** 2 years, 3 months ago

Definitely A. The CPU load for 1, 5, and 15 minutes is very high for executing/waiting processes.

upvoted 2 times

🗨️ 👤 **nixonbii** 2 years, 4 months ago

Definitely agree that A is the correct answer.

upvoted 1 times

🗨️ 👤 **KnifeClown1** 2 years, 4 months ago

Selected Answer: A

The correct answer is:

A. The system is under CPU pressure and will require additional vCPUs.

The load average (4.22 6.63 5.58) indicates that the system is under CPU pressure. The load average is the average number of tasks waiting for CPU time over a given time interval. If the load average is consistently high, it can mean that the system is struggling to keep up with demand and that more CPU resources are needed. In this case, the administrator may need to consider adding additional vCPUs to alleviate the CPU pressure.

upvoted 1 times

🗨️ 👤 **Nvoid** 2 years, 7 months ago

Selected Answer: B

real world, anyone would first reboot, and if it happened again frequently then upgrade. You don't go around changing things in IT and spending money to upgrade unless you have unlimited resources and/or money to throw around without any questions or consequences. You guys are thinking too hard on this one.

upvoted 4 times

🗨️ 👤 **Veteran903** 2 years, 7 months ago

nice try, but please explain how just rebooting the server is gonna help with load average?

upvoted 2 times

🗨️ 👤 **KnifeClown1** 2 years, 4 months ago

Option B is not accurate because the uptime of the system being up for over a year does not necessarily indicate that the system requires a reboot. Uptime is simply a measure of the time the system has been running without any downtime. It is possible for a system to run for much longer without requiring a reboot. Reboots are typically required to install updates or address specific performance or stability issues, but simply having a long uptime does not automatically imply a need for a reboot.


upvoted 1 times

A Linux administrator has set up a new DNS forwarder and is configuring all internal servers to use the new forwarder to look up external DNS requests. The administrator needs to modify the firewall on the server for the DNS forwarder to allow the internal servers to communicate to it and make the changes persistent between server reboots. Which of the following commands should be run on the DNS forwarder server to accomplish this task?

- A. `ufw allow out dns`
- B. `systemctl reload firewalld`
- C. `iptables -A OUTPUT -p udp -m udp -dport 53 -j ACCEPT`
- D. `firewall-cmd --zone=public --add-port=53/udp --permanent`

Suggested Answer: D

Community vote distribution



imnewtothis 9 months, 3 weeks ago

Selected Answer: D

This command adds a rule to the firewall using `firewall-cmd`, allowing UDP traffic on port 53, which is used for DNS communication. The `--permanent` option makes the change persistent between reboots.

upvoted 1 times

DRVision 1 year ago

Selected Answer: D

D. `firewall-cmd --zone=public --add-port=53/udp --permanent`. This command will open UDP port 53, which is the port used by DNS, in the public zone of the firewall. The `--permanent` option makes the change persistent across reboots.

Here's a brief explanation of the other options:

A. `ufw allow out dns` would be used if you were configuring a Ubuntu or Debian-based system that uses UFW ("Uncomplicated" Firewall - notice they did not specify uncomplicated which points towards a `firewalld`.)

B. `systemctl reload firewalld` is used to reload the firewall configuration, but it doesn't actually make any changes to the configuration.

C. `iptables -A OUTPUT -p udp -m udp -dport 53 -j ACCEPT` is an iptables command that would allow outgoing DNS traffic, but it wouldn't persist after a reboot unless saved to a file and restored on boot.

upvoted 3 times

Damon54 1 year, 4 months ago

Selected Answer: A

on Ubuntu

`ufw allow dns`

upvoted 1 times

Jacobmy98 1 year, 8 months ago

can `--permanent` be placed on the back end of the command ? ive always seen it on the front after `firewall-cmd`

upvoted 1 times

Jacobmy98 1 year, 8 months ago

turns out you can enter `--permanent` after. answer is D

upvoted 1 times

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

- A. chmod 775
- B. umask. 002
- C. chattr -Rv
- D. chown -cf

Suggested Answer: B

Community vote distribution



B (100%)

  **tutita** Highly Voted 1 year, 7 months ago

Selected Answer: B

keyword "by default" = umask

upvoted 5 times

  **ajna_** 1 year, 1 month ago

when you said this.

upvoted 3 times

  **bongobo** Most Recent 11 months, 2 weeks ago

correct answer is A chmod 775

upvoted 1 times

  **Ckl22** 2 years ago

Selected Answer: B

the umask octal value of 002 would take away the write permissions of "others"

upvoted 2 times

  **Jacobmy98** 1 year, 8 months ago

im confused. umask 002 allows write access to Other users. meaning anyone can modify or alter the file. Chmod 775 makes it so where the User and Group owner of the file is only able to modify while the 5 in Other makes it so they can only read and execute. NOT modify. Could anyone explain further if im wrong how its supposed to be ?

upvoted 1 times

  **examtopics11** 1 year, 7 months ago

B. umask. 002

default permissions Linux sets for files and directories before umask which is 666 for files and 777 for directories. 6-2 = 4 which is read only for Other users.

upvoted 5 times

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

- A. rsyslog
- B. cp
- C. rsync
- D. reposync
- E. scp
- F. ssh

Suggested Answer: CE

Community vote distribution

CE (100%)

  **abrilo** 6 months, 3 weeks ago

Selected Answer: CE

<https://www.golinuxcloud.com/commands-copy-file-from-one-server-to-another-linux-unix/>

upvoted 1 times

A Linux administrator rebooted a server. Users then reported some of their files were missing. After doing some troubleshooting, the administrator found one of the filesystems was missing. The filesystem was not listed in /etc/fstab and might have been mounted manually by someone prior to reboot. Which of the following would prevent this issue from reoccurring in the future?

- A. Sync the mount units.
- B. Mount the filesystem manually.
- C. Create a mount unit and enable it to be started at boot.
- D. Remount all the missing filesystems.

Suggested Answer: C

Community vote distribution

C (100%)

🗉 👤 **linux_admin** 10 months, 2 weeks ago

Selected Answer: C

C. Create a mount unit and enable it to be started at boot.

Explanation:

If a filesystem is missing after a reboot, it may be because it was not mounted at boot time. To prevent this from happening in the future, the filesystem should be mounted automatically at boot time. This can be done by creating a mount unit and enabling it to be started at boot.
upvoted 2 times

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

```
# grep joe /etc/passwd /etc/shadow
/etc/passwd:joe:x:1001:1001:./home/joe:/bin/nologin
/etc/shadow:joe:$6$3uOw6qWx9876jGhgKJsdFH987634534voj.:18883:0:99999:7:::
```

Which of the following commands would resolve the issue?

- A. `usermod -s /bin/bash joe`
- B. `pam_tally2 -u joe -r`
- C. `passwd -u joe`
- D. `chage -E 90 joe`

Suggested Answer: A

Community vote distribution

A (100%)

POGActual 9 months, 1 week ago

Answer is A. The current login shell is set to `/bin/nologin` which, you can guess, doesn't allow the user to login when this is set. The login shell needs changed to `/bin/bash` to allow login.

upvoted 2 times

linux_admin 10 months, 2 weeks ago

Selected Answer: A

The command `usermod -s /bin/bash joe` changes the login shell of the user named "joe" to `/bin/bash`.

Explanation:

`usermod` is a command in Linux used to modify user account details.

`-s /bin/bash` is an option used with the `usermod` command to specify the new login shell for the user. In this case, the login shell is changed to `/bin/bash`.

"joe" is the username of the user whose account is being modified.

Changing the login shell of a user can be useful in cases where a different shell is required for a particular user. In this case, the user "joe" will now use the Bash shell as their default shell when they log in.

upvoted 4 times

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

- A. `rpm -qf /etc/httpd/conf/httpd.conf`
- B. `rpm -ql /etc/httpd/conf/httpd.conf`
- C. `rpm --query /etc/httpd/conf/httpd.conf`
- D. `rpm -q /etc/httpd/conf/httpd.conf`

Suggested Answer: A

Community vote distribution

A (100%)

🗉 **linux_admin** 10 months, 2 weeks ago

Selected Answer: A

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

Explanation:

`rpm` is a command-line package manager used in many Linux distributions, including Red Hat, CentOS, and Fedora.

`-qf` is an option used with the `rpm` command to query the package that owns a particular file.

`/etc/httpd/conf/httpd.conf` is the path of the configuration file that is missing.

So, the command `rpm -qf /etc/httpd/conf/httpd.conf` will return the name of the RPM package that installed the configuration file on the system
upvoted 3 times

Users have been unable to save documents to /home/tmp/temp and have been receiving the following error:

Path not found -

A junior technician checks the locations and sees that /home/tmp/tempa was accidentally created instead of /home/tmp/temp. Which of the following commands should the technician use to fix this issue?

- A. cp /home/tmp/tempa /home/tmp/temp
- B. mv /home/tmp/tempa /home/tmp/temp
- C. cd /tmp/tmp/tempa
- D. ls /home/tmp/tempa

Suggested Answer: B

Community vote distribution

B (100%)

🗨️ 👤 **linux_admin** 10 months, 2 weeks ago

Selected Answer: B

B. mv /home/tmp/tempa /home/tmp/temp

Explanation:

mv is a command in Linux used to move or rename files and directories.

/home/tmp/tempa is the path of the accidentally created directory.

/home/tmp/temp is the correct path of the directory where users should be saving their documents.

The command mv /home/tmp/tempa /home/tmp/temp will move the directory /home/tmp/tempa to /home/tmp/temp, effectively correcting the issue.

upvoted 3 times

🗨️ 👤 **Ckl22** 1 year ago

Selected Answer: B

B is correct

mv old-filename new-filename

upvoted 2 times

A Linux administrator has logged in to a server for the first time and needs to know which services are allowed through the firewall. Which of the following options will return the results for which the administrator is looking?

- A. firewall-cmd --get-services
- B. firewall-cmd --check-config
- C. firewall-cmd --list-services
- D. systemctl status firewalld

Suggested Answer: C

Community vote distribution

C (100%)

🗉 👤 **linux_admin** 10 months, 2 weeks ago

Selected Answer: C

C. firewall-cmd --list-services

Explanation:

firewall-cmd is the command-line interface for the firewall management tool firewalld in Linux.

--list-services is an option used with the firewall-cmd command to list all the services that are currently allowed through the firewall.

So, the command firewall-cmd --list-services will return the list of services that are currently allowed through the firewall.

upvoted 2 times

After installing a new version of a package, a systems administrator notices a new version of the corresponding .service file was installed. In order to use the new version of the .service file, which of the following commands must be issued FIRST?

- A. systemctl status
- B. systemctl stop
- C. systemctl reinstall
- D. systemctl daemon-reload

Suggested Answer: D

Community vote distribution

D (100%)

linux_admin 10 months, 2 weeks ago

Selected Answer: D

D. systemctl daemon-reload

Explanation:

systemctl is a command-line tool used to manage services in Linux.

daemon-reload is an option used with the systemctl command to reload the systemd manager configuration. This is necessary when changes are made to the configuration files of a service, such as the .service file.

After a new version of a package is installed, the corresponding .service file may be updated. In order to use the new version of the .service file, it is necessary to reload the configuration with the systemctl daemon-reload command before any other action can be taken. This ensures that the systemd manager is aware of the changes made to the .service file.

upvoted 2 times

Ckl22 1 year ago

Selected Answer: D

D is correct

To have your changes take effect, issue the systemctl daemon-reload command for the service whose unit file you modified or extended. After you accomplish that task, issue the systemctl restart command to start or restart the service.

upvoted 1 times

Which of the following commands is used to configure the default permissions for new files?

- A. setenforce
- B. sudo
- C. umask
- D. chmod

Suggested Answer: C

Community vote distribution

C (100%)

linux_admin 1 year, 4 months ago

Selected Answer: C

C. umask

Explanation:

umask is a command in Linux that sets the default file permissions for newly created files and directories.

The default permissions are calculated by subtracting the umask value from 777 for files and 666 for directories. The result is the default permission for each bit (r, w, and x).

For example, a umask value of 022 would result in file permissions of 644 (666 - 022) and directory permissions of 755 (777 - 022).

Option A (setenforce) is used to set the enforcement mode of SELinux, the security module in Linux. It is not used to configure the default permissions for new files.

Option B (sudo) is used to run commands with administrative privileges. It is not used to configure the default permissions for new files.

Option D (chmod) is used to change the permissions of existing files and directories, but not to configure the default permissions for new files.
upvoted 2 times

Damon54 10 months, 1 week ago

The default permissions are calculated by subtracting the umask value from 777 for Directories and 666 for files

upvoted 1 times

Ckl22 1 year, 6 months ago

Selected Answer: C

C is correct

The user mask (umask) feature defines the default permissions Linux assigns to the file or directory. The user mask is an octal value that represents the bits to be removed from the octal mode 666 permissions for files or the octal mode 777 permissions for directories.

upvoted 2 times

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

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

Suggested Answer: D

Community vote distribution

D (100%)

🗨️ 👤 **nixonbii** 10 months, 1 week ago

Are we even going to address the implications of this question? Never, ever enable remote root login. I understand the concept, but the implied methodology is very scary.

upvoted 3 times

🗨️ 👤 **linux_admin** 10 months, 2 weeks ago

Selected Answer: D

D. `ssh -i "ABC-key.pem" root@10.0.0.1`

Explanation:

ssh is a command in Linux used to establish a secure shell connection to a remote server.

-i "ABC-key.pem" is an option used with the ssh command to specify the private key file to use for authentication. The private key file is used to establish a secure connection to the remote server.

root@10.0.0.1 is the username and IP address of the remote server to which the connection is being established.

upvoted 2 times

In order to copy data from another VLAN, a systems administrator wants to temporarily assign IP address 10.0.6.5/24 to the newly added network interface enpls0f1. Which of the following commands should the administrator run to achieve the goal?

- A. `ip addr add 10.0.6.5/24 dev enpls0f1`
- B. `echo "IPV4_ADDRESS=10.0.6.5/24" > /etc/sysconfig/network-scripts/ifcfg-enpls0f1`
- C. `ifconfig 10.0.6.5/24 enpls0f1`
- D. `nmcli conn add ipv4.address=10.0.6.5/24 if name snpls0f1`

Suggested Answer: A*Community vote distribution*A (100%)

🗳️ **bongobo** 9 months, 2 weeks ago

if ifconfig is involved, shout be:

```
ifconfig 10.0.6.5 netmask 255.255.255.0 -I enpls0f1 up
```

upvoted 2 times

🗳️ **bongobo** 9 months, 2 weeks ago

```
ifconfig -I enpls0f1 10.0.6.5 netmask 255.255.255.0 up
```

upvoted 2 times

🗳️ **linux_admin** 1 year, 10 months ago

Selected Answer: A

A. `ip addr add 10.0.6.5/24 dev enpls0f1`

Explanation:

`ip` is a command-line tool used to manage network interfaces in Linux.

`addr add 10.0.6.5/24` is an option used with the `ip` command to add an IP address to a network interface.

`dev enpls0f1` specifies the network interface to which the IP address should be added.

So, the command `ip addr add 10.0.6.5/24 dev enpls0f1` will assign IP address 10.0.6.5/24 temporarily to the network interface enpls0f1.

upvoted 3 times

A Linux engineer needs to download a ZIP file and wants to set the nice value to -10 for this new process. Which of the following commands will help to accomplish the task?

- A. `$ nice -v -10 wget https://foo.com/installation.zip`
- B. `$ renice -v -10 wget https://foo.com/installation.zip`
- C. `$ renice -10 wget https://foo.com/installation.zip`
- D. `$ nice -10 wget https://foo.com/installation.zip`

Suggested Answer: D

Community vote distribution

D (100%)

🗨️ **Damon54** 10 months, 1 week ago

Selected Answer: D

although the -n option should be specified

`nice -n 10 wget https://foo.com/installation.zip`

upvoted 3 times

🗨️ **linux_admin** 1 year, 4 months ago

Selected Answer: D

D. `$ nice -10 wget https://foo.com/installation.zip`

Explanation:

`nice` is a command in Linux used to run a command with a modified scheduling priority.

`-10` is an option used with the `nice` command to set the nice value to -10, which increases the priority of the command.

`wget` is a command in Linux used to download files from the internet.

`https://foo.com/installation.zip` is the URL of the ZIP file to be downloaded.

So, the command `nice -10 wget https://foo.com/installation.zip` will download the ZIP file and set the nice value of the `wget` process to -10, which will increase its priority.

upvoted 3 times

Which of the following data structures is written in JSON?

-
- A. `name: user1`
`position: DevOps`
`floor: 3`
- B. `<table>`
`<tbody><tr>`
`<td>user1</td>`
`<td>DevOps</td>`
`<td>3</td>`
`</tr>`
`</tbody></table>`
- C. `<root>`
`<floor>3</floor>`
`<name>user1</name>`
`<position>DevOps</position>`
`</root>`
- D. `{`
 `"name": "user1",`
 `"job": "DevOps",`
 `"floor": 3`
`}`

Suggested Answer: D

Community vote distribution

D (100%)

 **linux_admin** 10 months, 2 weeks ago

Selected Answer: D

The syntax of JSON is based on a collection of key-value pairs that are separated by commas and enclosed in curly braces { }. The keys are always strings, enclosed in double quotes, followed by a colon, and then the value.

Values in JSON can be of several types, including strings, numbers, booleans, null, arrays, and objects. String values are enclosed in double quotes, numeric values do not need quotes, and boolean values are true or false. Null values are represented as the null keyword.

Arrays are a collection of values, enclosed in square brackets [], and separated by commas. Objects are a collection of key-value pairs, enclosed in curly braces { }, and separated by commas. Objects can be nested within other objects, and arrays can contain other arrays or objects.

upvoted 3 times

Which of the following enables administrators to configure and enforce MFA on a Linux system?

- A. Kerberos
- B. SELinux
- C. PAM
- D. PKI

Suggested Answer: A

Community vote distribution

C (100%)

🗳️ **nixonbii** 10 months, 1 week ago

C is correct, but so is A. Kerberos uses the ticket granting ticket system often used for federated trust relationships. For each resources requiring additional verification, a Kerberos system can require additional information such as a token or a verification code sent to a known phone number. Since this is a Linux exam, PAM will suffice.

upvoted 1 times

🗳️ **linux_admin** 10 months, 2 weeks ago

Selected Answer: C

The technology that enables administrators to configure and enforce Multi-Factor Authentication (MFA) on a Linux system is:

C. PAM (Pluggable Authentication Modules)

Explanation:

PAM (Pluggable Authentication Modules) is a framework that allows the integration of multiple authentication schemes for Linux systems. It provides a way for administrators to configure and enforce MFA for user authentication by enabling the use of multiple factors, such as passwords, smart cards, biometrics, and others.

PAM allows the configuration of authentication policies for different services and applications running on a Linux system. This makes it possible to implement MFA for specific applications or services, rather than applying it to the entire system. Additionally, PAM allows administrators to configure different authentication methods for different users or user groups, depending on their access requirements.

upvoted 4 times

🗳️ **linux_admin** 10 months, 2 weeks ago

Kerberos is a network authentication protocol used for single sign-on (SSO), but it does not provide MFA capabilities.

upvoted 2 times

🗳️ **KnifeClown1** 10 months, 3 weeks ago

Selected Answer: C

C. PAM (Pluggable Authentication Modules) enables administrators to configure and enforce multi-factor authentication (MFA) on a Linux system. PAM is a modular architecture that allows system administrators to specify how authentication should be performed on their systems, including the use of MFA. This can be achieved by adding additional modules to the PAM configuration, such as modules for OTPs, smart cards, or biometric authentication. PAM can also be used to integrate with external authentication systems, such as LDAP, Kerberos, or RADIUS, to enforce MFA for authentication to various services on the Linux system.

upvoted 2 times

🗳️ **Nvoid** 1 year, 1 month ago

Selected Answer: C

Choosing C here.


upvoted 3 times

🗳️ **SaadiaS** 1 year, 1 month ago

C is correct. Pluggable Authentication Modules allow Linux to work with Google Authenticator and other OTP tools to add two-factor security to your system.

<https://www.redhat.com/sysadmin/mfa-linux>

upvoted 3 times

 **MissAllen** 1 year, 1 month ago

Pluggable Authentication Modules (PAM) provide for multi-factor authentication on Linux. Answer C.

upvoted 4 times

A database administrator requested the installation of a custom database on one of the servers. Which of the following should the Linux administrator configure so the requested packages can be installed?

- A. /etc/yum.conf
- B. /etc/ssh/sshd.conf
- C. /etc/yum.repos.d/db.repo
- D. /etc/resolv.conf

Suggested Answer: C

Community vote distribution

C (100%)

🗉 👤 **linux_admin** 10 months, 2 weeks ago

Selected Answer: C

C. /etc/yum.repos.d/db.repo

Explanation:

YUM (Yellowdog Updater, Modified) is a package manager that allows installation, removal, and update of software packages on Linux systems. YUM uses repositories (usually hosted on the internet) to retrieve the necessary packages.

The /etc/yum.repos.d directory contains the configuration files for the repositories used by YUM. By default, YUM is configured to use the repositories listed in the /etc/yum.conf file, but custom repositories can be added to the /etc/yum.repos.d directory.

upvoted 4 times

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

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

Suggested Answer: A

Community vote distribution

A (100%)

 **linux_admin** Highly Voted 10 months, 2 weeks ago

Selected Answer: A

A. /var

Explanation:

In Linux, system logs are commonly stored in the /var (variable) filesystem, specifically in the /var/log directory. This directory contains logs for various system services and applications, such as the syslog, mail, and authentication logs.

The /tmp directory is a temporary filesystem that is used to store temporary files and directories created by system processes and applications. It is typically cleared on system reboot, and its contents should not be relied on for long-term storage.

The /etc directory contains configuration files for the system and various applications, but it does not typically store system logs.

The /opt directory is used to install additional software packages, but it is not a standard location for storing system logs.

upvoted 5 times

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

- A. kill -1
- B. kill -3
- C. kill -15
- D. kill -HUP
- E. kill -TERM

Suggested Answer: B

Community vote distribution

C (100%)

🗨️ **TheRealManish** Highly Voted 2 years, 7 months ago

It seems to me that the correct answer is not here.

Kill -9 is the right answer, all of the other answers are wrong.

Kill -1 restart process

kill -3 send a control c that can be ignored

kill -15 = request a graceful shutdown that can be ignored

kill -HUP same as kill -1

kill -TERM same as kill -15

upvoted 13 times

🗨️ **diabolee** Most Recent 6 months, 4 weeks ago

Selected Answer: B

Going with B since kill -15 gracefully shuts down the process.

upvoted 1 times

🗨️ **Blatzzy** 1 year, 6 months ago

To forcibly stop a process, the administrator should use the -9 signal, also known as SIGKILL. Therefore, the correct option is not listed among the provided choices. If you have the option to provide a custom signal, you can use kill -9 followed by the process ID.

In this case, if you must choose from the given options:

B. kill -3 sends a SIGQUIT signal, which can also terminate a process, but it allows the process to perform a core dump before exiting. It's a bit less forceful than SIGKILL.

So, the closest option among the given choices is B. kill -3. However, if the goal is to forcefully terminate the process, using kill -9 would be more appropriate.

upvoted 3 times

🗨️ **salthedhash** 1 year, 7 months ago

Selected Answer: C

The -15 signal (SIGTERM) is used to request a process to terminate. It allows the process to perform cleanup operations before exiting. If a process doesn't respond to SIGTERM, more aggressive signals like -9 (SIGKILL) can be used, but it does not allow the process to perform any cleanup. The administrator might resort to using -9 only if a process is unresponsive to the regular termination signals.

upvoted 1 times

🗨️ **linux_admin** 2 years, 4 months ago

A. kill -1 (or kill -SIGHUP) - Sends a SIGHUP signal to the specified process, which typically requests the process to reload its configuration files. This may not terminate the process, so it may not be the best choice for forcibly stopping a process.

B. kill -3 (or kill -SIGQUIT) - Sends a SIGQUIT signal to the specified process, which requests the process to terminate and create a core dump. This may not terminate the process, so it may not be the best choice for forcibly stopping a process.

C. kill -15 (or kill -SIGTERM) - Sends a SIGTERM signal to the specified process, which requests the process to terminate gracefully. This is the default signal and is the recommended way to stop a process, as it allows the process to clean up its resources before exiting. However, some processes may not respond to the SIGTERM signal, in which case a more forceful approach may be necessary.

upvoted 2 times

🗨️ 👤 **linux_admin** 2 years, 4 months ago

D. kill -HUP - This is the same as kill -1, which may not be the best choice for forcibly stopping a process.

E. kill -TERM - This is the same as kill -15, which may not be the best choice for forcibly stopping a process.

If the process needs to be forcibly stopped, the best option would be to use kill -9 (or kill -SIGKILL), which sends a SIGKILL signal to the specified process. This signal immediately terminates the process without giving it a chance to clean up its resources. However, this should only be used as a last resort, after other options (such as SIGTERM) have been tried and failed.

upvoted 1 times

🗨️ 👤 **MaryamNesa** 2 years, 7 months ago

C is correct

upvoted 1 times