Linux
Interview Questions

Click here to view the live version of the page

Top Answers to Linux Interview Questions

Linux is among the fastest and most powerful operating systems used in computers. Over 90 percent of the world's fastest computers have Linux OS. If you wish to become a Linux professional in a reputed organization, then this is the right platform to prepare for your job interview. In this Linux Interview Questions blog, you will cover some of the most common interview questions asked during interviews in this domain. Let's get a quick look at these frequently asked questions:

We have categorized the Linux interview questions and answers these parts as mentioned below:

Basic Linux Interview Questions for Freshers

Advanced Linux Interview Questions for Experienced

Linux Interview Questions for 2 to 3 Years of Experience

Linux Interview Questions for 4, 5 Years and More Experience

Linux Commands Interview Questions

Linux Admin Interview Questions

Linux Troubleshooting Interview Questions

Linux Networking Interview Questions

Linux Professionals' Salary Trends

Linux Professionals' Job Trends

Linux Professionals' Roles & Responsibilities

Conclusion

Did you Know?

- According to statistics of the top 1 million web servers, 96.3% are running on Linux.
- According to reports from GlobeNewswire, by 2027, the Linux operating system market is projected to achieve a value of $15.64 billion.
- In the github repository of the Linux kernel, there are around 27.8 million lines of code.

Basic Linux Interview Questions for Freshers

1. What is Linux?

Linux is an open-source operating system based on UNIX. It was named after the founder “Linus Torvalds”. He introduced Linux with the primary goal to offer an operating system at a free or very reasonable price for users. It is based on the Linux kernel and is compatible with different hardware platforms such as Intel, MIPS, HP, IBM, SPARC, and Motorola hardware platforms. Linux’s mascot, a penguin named Tux, is another popular feature. Linux offers a user-friendly environment where they can easily modify and create variations in the source code.

2. Compare Linux with Windows.

Criteria	Linux	Windows
Type of OS	Open-source	Proprietary
Customization	High level of customization	Cannot be changed

Security	Excellent	Vulnerable to security issues

3. What are the components of the Linux system?

There are three primary components of the Linux system which are explained below.

Kernel: The kernel is the most important component of Linux. It is in charge of the operating system's primary functions. It is made up of a number of modules that interface directly with the hardware. Kernel offers the necessary abstraction for system or application programs to mask low-level hardware information.

System libraries: They are specialized functions or programs that allow application programs or system utilities to access Kernel capabilities. These libraries implement the majority of the operating system's functionality and do not require kernel module code access permissions.

System Utility: Programs in the System Utility category are in charge of performing specialized, individual-level activities. They are more dependable and also provide users control over the computer.

4. What is LILO?

LILO (Linux Loader) is a boot loader for Linux. It is used to load Linux into memory and start the operating system. LILO can be configured to boot other operating systems as well. LILO is customizable, i.e., if the default configuration is not correct, it can be changed. lilo.conf is the configuration file for LILO. LILO is also a code snippet that loads PC BIOS into the main memory at the time of starting the computer system.

It handles the following tasks:

- Locating Linux kernel

- Identifying other supporting programs and loading them in memory
- Starting the kernel

The selection of various Kernel images and boot routines is supported by LILO. For this reason, it is known as the boot manager.

5. Suppose, you wish to print a file 'draft' with 60 lines on a page. What command would you use?

The command used for this purpose would be as follows:

```
pr -l60 draft
```

Note: The default page length when using pr is 66 lines. The -l option specifies a different length.

6. What is LD_LIBRARY_PATH?

LD_LIBRARY_PATH is an environment variable used for debugging a new library or a non-standard library. It is also used to identify the directories that need to be searched for; in order to do this, the path to search for the directories needs to be specified.

The variable can be set using the following:

```
setenv—LD_LIBRARY_PATH--$PATH
```

It is used to search for the shared objects/dynamic libraries by the operating system for extendable functionality at the runtime.

Prepare yourself for the Linux certification with this comprehensive Linux Training in London!

7. Name a service that you should disable (which acts both as web and FTP servers) on a Linux server.

The finger service should be disabled on a Linux server because a remote user can get important information about the system by using this command.

To disable this service use command: sudo systemctl disable vsftpd

8. What does sar provide? Where are the sar logs stored?

The sar command in Linux is a valuable tool for collecting and analyzing system activity information. It reports various aspects of system performance, such as CPU usage, memory utilization, disk activity, network traffic, and more. When troubleshooting performance issues, sar enables you to review historical data and identify the causes of high load on specific system components.

When the CPU utilization is close to 100 percent, it indicates that the processing workload primarily demands the CPU. This information helps determine if the system is experiencing a CPU-bound situation.

By default, sar saves its log files in the /var/log/sa/sadd directory, where “dd” represents the current day. These log files are valuable for retrospective analysis and tracking system activity over time.

9. How to check memory stats and CPU stats as a Linux Admin?

Using the free and vmstat commands, we can display the physical and virtual memory statistics, respectively. With the help of the sar command, we can see the CPU utilization and other stats.

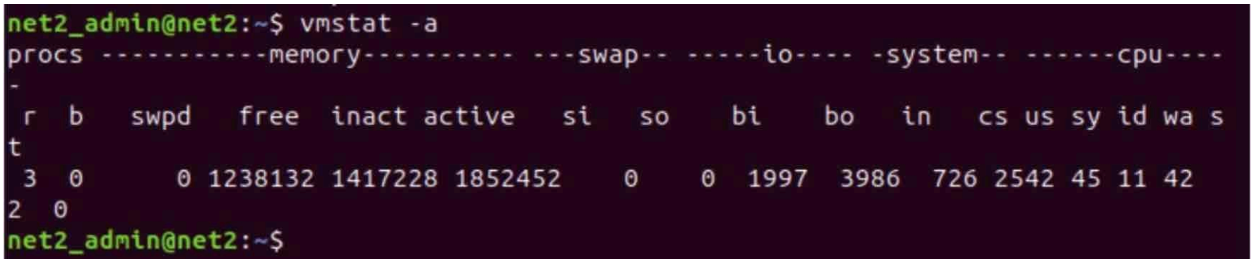

```
net2_admin@net2:~$ vmstat -a
procs -----------memory---------- ---swap-- -----io---- -system-- ------cpu----
-
 r  b   swpd   free  inact active   si   so    bi    bo   in   cs us sy id wa s
t
 3  0      0 1238132 1417228 1852452    0    0  1997  3986  726 2542 45 11 42
2  0
net2_admin@net2:~$
```

10. How to reduce or shrink the size of the LVM partition?

Below are the logical steps to reduce the size of the LVM partition:

- Unmount the file system using the unmount command
- Use the resize2fs command as follows:

```
resize2fs /dev/mapper/myvg-mylv 10G
```

- Then, use the lvreduce command as follows:

```
lvreduce -L 10G /dev/mapper/myvg-mylv
```

This way, we can reduce the size of the LVM partition and fix the size of the file system to 10 GB.

11. What are the different modes of Network Bonding in Linux?

Below is the list of various modes used in Network Bonding:

- balance-rr or mode 0: The round-robin mode for fault tolerance and load balancing
- active-backup or mode 1: Sets the active-backup mode for fault tolerance

- balance-xor or mode 2: Sets an XOR (exclusive-or) mode for fault tolerance and load balancing
- broadcast or mode 3: Sets a broadcast mode for fault tolerance. All transmissions are sent on all the slave interfaces
- 802.3ad or mode 4: Sets an IEEE 802.3ad dynamic link aggregation mode and creates aggregation groups that share the same speed and duplex settings
- balance-tlb or mode 5: Sets a transmit load balancing (TLB) mode for fault tolerance and load balancing
- balance-alb or mode 6: Sets an active load balancing (ALB) mode for fault tolerance and load balancing

For more details, check out Intellipaat's Linux Training in Sydney!

12. How to check and verify the status of the bond interface?

Using the following command, we can check which mode is enabled and what LAN cards are used in this bond:

```
cat /proc/net/bonding/bond0
```

In this example, we have a single bond interface. However, we can have multiple bond interfaces like bond1, bond2, and so on.

13. Do you know the Maximum length (in bytes) of the filename in Linux?

The maximum length of a filename is 255 bytes. In this filename, the pathname is not included, so the total length of the pathname and filename may easily surpass 255 characters.

14. What are the two different kinds of Linux User Modes?

The following are the two types of Linux user modes:

- Command Line
- GUI

15. What is Hard Link?

In Linux, Hard links can be defined as another name for an already existing file. For each file, we can generate an unlimited number of hard links. They have the ability to generate links for other hard connections. We can use the `Is-I` command to find out the total number of hard links in a file. And we can create Hard links using the following command:

```
$ ln [original filename] [link name]
```

16. What is Soft Link?

Soft link is also known as a symbolic link. Soft links are files that, in most cases, lead to another file. It just links to another entry somewhere in the file system and does not include any data in the destination file. These kinds of connections can be utilized across several file systems. The following command can be used to create soft links:

```
$ ln -s [original filename] [link name]
```

Check out this video on Linux Shell Tutorial:

17. What is a shell in Linux?

A shell is a command-line interface that allows users to interact with the Linux operating system. It acts as an interpreter between the user and the kernel by executing commands and returning output. Some common Linux shells are Bash, Csh, Ksh, and Zsh, among others. These provide an environment to automate tasks through scripting and access system resources.

18. What are daemons in Linux?

Daemons are background processes that start at system boot and keep running to perform system-critical tasks. Some examples are httpd daemon, which runs the Apache web server, and sshd daemon, which handles SSH remote connections. Daemons have no controlling terminal and run in the background without user intervention. They can be controlled using init scripts to start, stop, and check status.

19. What is the difference between an absolute path and a relative path in Linux?

An absolute path specifies the full location of a file or directory from the root directory (/), for example, /usr/local/bin/python. A relative path provides the location relative to the current working directory. For example, if the current directory is /usr/local, the relative path to the Python file is bin/python.

20. How are environment variables and shell variables different in Linux?

Environment variables are available to all child processes spawned from that shell or environment. Shell variables are internal variables within that particular shell. Environment variables use the export command, while shell variables do not.

21. What are the key benefits of scripting on Linux?

Some benefits of scripting on Linux include automating repetitive tasks, bulk administration of systems, error-free execution, cross-platform functionality, and reusability of code. Scripting saves time and effort compared to performing tasks manually. It also leads to standardized processes and less human error.

22. What is grep command, and how is it useful on Linux?

Grep is used to search for specified patterns within text files or command outputs. It can search for plain text, regular expressions, and match multiple files. This allows for quickly filtering large data sources to find the required info. Useful options include -i for case-insensitive search, -R to recursively search directories, and –color to highlight matches.

23. What is an array in Linux shell scripting? Give an example.

Arrays allow the storage of multiple data elements into a single variable. It stores elements as indexed values and can be iterated through loop statements.

For example,

```
fav_fruits=(Apple Mango Banana)

echo ${fav_fruits[0]} # Prints Apple

for fruit in ${fav_fruits[@]}; do

    echo $fruit

done
```

24. How is memory management handled in Linux?

Linux manages memory efficiently with advanced techniques like virtual memory, shared libraries, demand paging, and swap space. It allocates memory to processes only when needed and shares common libraries across multiple processes. Virtual memory maps some physical memory onto disk for overflow, while demand paging loads executable code pages on demand.

25. How do you manage software packages and libraries in Linux?

Linux distros include dedicated package manager tools like apt, yum, rpm, etc., to find, install, remove, and update packages easily. Packages contain compiled binaries, libraries, docs, and configs packaged neatly for dependency resolution. Popular frameworks like rpm and deb have repositories to fetch both official and third-party packages.

26. What is a Linux distribution? Name some popular Linux distributions.

A Linux distribution consists of the Linux kernel along with collections of software bundled together and optimized to run on top of the kernel. Some popular general-purpose distros are Ubuntu, Fedora, Debian, openSUSE, etc. Specialized distros also exist, like Kali Linux for security and CentOS for enterprise servers, among others.

Advanced Linux Interview Questions for Experienced

27. Explain the features of the Linux system?

The key features of the Linux system are as follows:

- Linux is a community-based major project which is freely available open-source code. Multiple teams collaborate to improve the capabilities of this operating system, which is always growing.
- It offers a prominent feature which is that it is a multiuser system, which implies that several users may share system resources such as memory, ram, and application programs.
- Portability refers to the capacity of software to run on a variety of hardware platforms in the same way. The Linux kernel and application software may be installed on virtually any hardware platform.
- Linux is a multiprogramming system, which means it can run many programs at the same time.
- Linux has a Hierarchical File System (HFS), which offers a standardized structure for storing system and user data files.
- Linux contains a custom interpreter application that allows users to run operating system program commands and instructions.
- User security is provided by Linux through authentication mechanisms such as password protection, limited access to particular files, and data encryption.

Do checkout our blog on top features of linux operating system to gain in-depth knowledge about it!

28. Explain various file permissions in Linux?

In Linux, each file and directory has three categories of owners which are User, Group, and Others. For all three owners, there are three sorts of permissions defined as mentioned below:

Read: This read permission allows you to open the file, read it, and list the content of the directory.

Write: This permission gives you the ability to change the contents of a file as well as add, remove, and rename files in directories.

Execute: The file in the directory can be accessed and run by the user. The execute permission must be established before a file may be run.

29. Why is Linux regarded as a more secure operating system than other operating systems?

Linux has become more popular in the technology industry in terms of security. There are several reasons why Linux is more secure than other operating systems.

- On Linux, only a few people have access to the system. As a result, the virus cannot infect the entire system but it may affect only a few files.
- Before opening the files, Linux users must first complete the tasks, so that they can protect their systems against flaws.
- The Linux operating system includes a variety of working environments, including Linux Mint, Debian, Arch, and others, all of which include virus protection.
- It keeps a log history so that it may quickly see the specifics of the system files afterward.
- Iptables is a Linux feature that examines the system’s security circle.
- As Linux users are comparatively fewer in number as compared to other operating systems, security will be enhanced.

Learn Linux from Top Experts by Enrolling in Linux Training in Dubai.

30. Which command is used to check the number of files, disk space, and each user’s defined quota?

The repquota command is used to check the status of a user’s defined quota, along with the disk space and the number of files used.

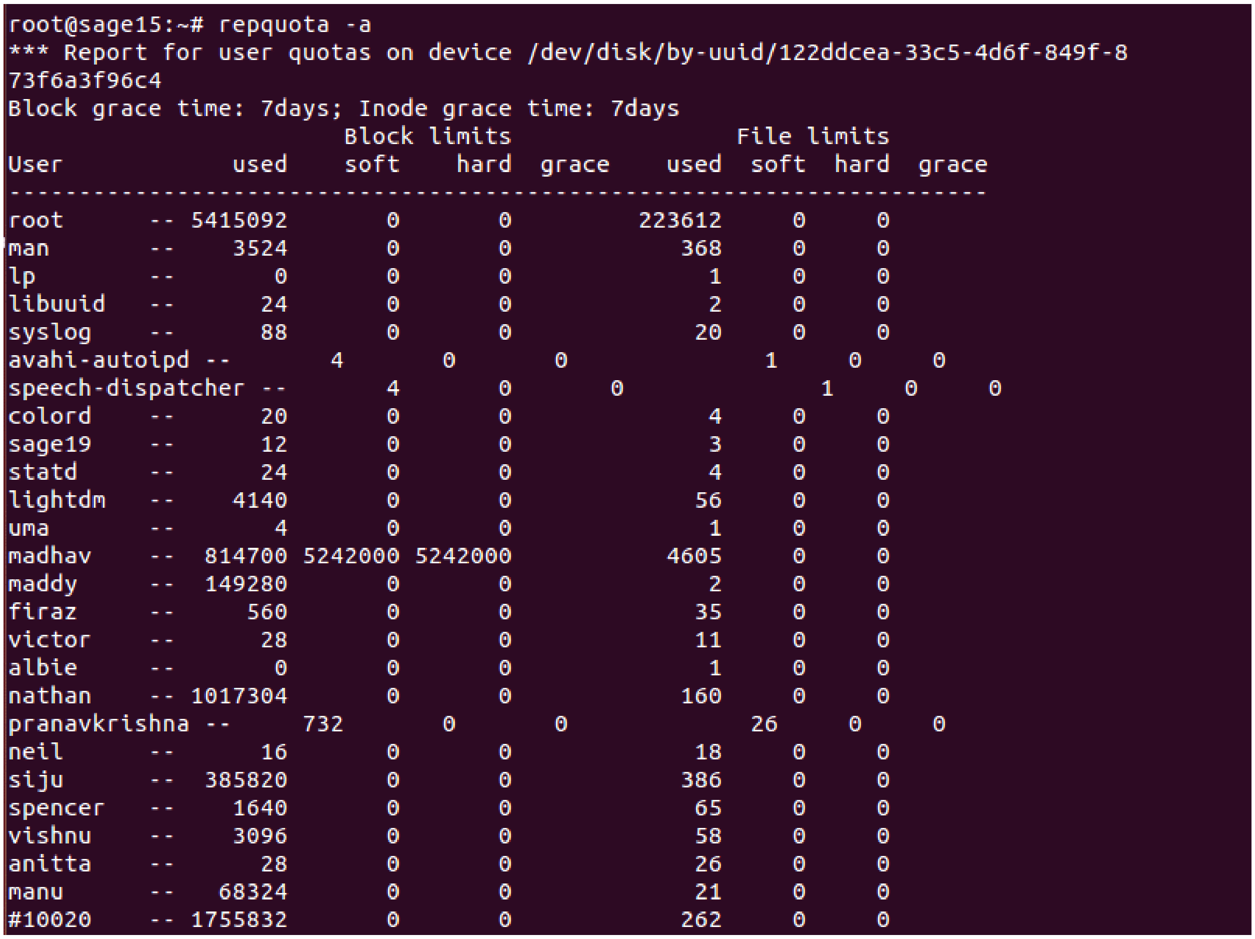

```
root@sage15:~# repquota -a
*** Report for user quotas on device /dev/disk/by-uuid/122ddcea-33c5-4d6f-849f-8
73f6a3f96c4
Block grace time: 7days; Inode grace time: 7days
                        Block limits                File limits
User            used    soft    hard  grace    used  soft  hard  grace
----------------------------------------------------------------------
root      -- 5415092       0       0         223612     0     0
man       --    3524       0       0            368     0     0
lp        --       0       0       0              1     0     0
libuuid   --      24       0       0              2     0     0
syslog    --      88       0       0             20     0     0
avahi-autoipd --       4       0       0              1     0     0
speech-dispatcher --       4       0       0              1     0     0
colord    --      20       0       0              4     0     0
sage19    --      12       0       0              3     0     0
statd     --      24       0       0              4     0     0
lightdm   --    4140       0       0             56     0     0
uma       --       4       0       0              1     0     0
madhav    --  814700 5242000 5242000           4605     0     0
maddy     --  149280       0       0              2     0     0
firaz     --     560       0       0             35     0     0
victor    --      28       0       0             11     0     0
albie     --       0       0       0              1     0     0
nathan    -- 1017304       0       0            160     0     0
pranavkrishna --     732       0       0             26     0     0
neil      --      16       0       0             18     0     0
siju      --  385820       0       0            386     0     0
spencer   --    1640       0       0             65     0     0
vishnu    --    3096       0       0             58     0     0
anitta    --      28       0       0             26     0     0
manu      --   68324       0       0             21     0     0
#10020    -- 1755832       0       0            262     0     0
```

This command gives a summary of the user's quota, i.e., how much space and files are left for the particular user. Each user has a defined quota in Linux. This is done mainly for security as it restricts files from unwanted access. The quota can be given to a single user or to a group of users.

Linux Interview Questions for 2 to 3 Years of Experience

31. How can you enhance the security of the password file in Linux?

It is in the test file named '/etc/passwd' that Linux usually keeps its user account details, including the one-way encrypted passwords. However, this file can be accessed with the help of different tools, which might throw security issues.

To minimize this risk, we will make use of the shadow password format that saves the account details in a regular file /etc/passwd as in the traditional method but with the password stored as a single 'x' character, i.e., it is not the original password that is actually stored in this file. Meanwhile, a second file /etc/shadow will have the encrypted password, along with the other relevant information, such as the account/password expiration date, etc. Most importantly, the latter file is readable only by the root account, and thus it minimizes the security risk.

To enable shadow password use the command: pwconv

32. What are the three standard streams in Linux?

In Linux, standard streams are channel communication of input and output between a program and its environment. In the Linux system, input and output are spread among three standard streams which are:

1. Standard Input (stdin)
2. Standard Output (stdout)
3. Standard Error (stderr)

33. What command can you use to make a tape archive file of /home directory and send it to the /dev/tape device?

The command used here is:

```
tar -cvf /dev/tape /home
```

The -xvf option is used to extract files from an archive.

Expert Linux professionals are in high demand. Take this Linux Course in Toronto and join the big league!

34. What is CLI?

The acronym CLI stands for Command Line Interface. The user can input declarative instructions into this interface to direct the machine tasks. It communicates with a software program by issuing commands in the form of text lines. It also interacts with computer terminals; the interface receives text lines and transforms them into operating system commands. CLI offers great flexibility.

35. What is GUI?

A GUI (Graphical User Interface) is a type of interface between humans and machines that allows people to interact with electronic devices via graphical icons and visual indications. The inclusion of graphical components makes it easier to interact with the system, as well as provides additional appeal through images, icons, and colors, rather than having to memorize and write commands. Users will find it simpler to engage with the system if certain graphical components or icons are used. It is visually appealing and enables increased productivity.

36. Suppose, your FTP Server is hacked and the entire server needs to be restored. How would you restore the original kernel system files?

We cannot restore the entire operating system from the tape backup device. Therefore, we should reinstall the core operating system and then restore the system configuration files and user data from the tape backup device.

37. Why should you avoid Telnet to administer a Linux system remotely?

Telnet uses the most insecure method for communication. It sends data across the network in plain text format, and anybody can easily find out the password using the network tool.

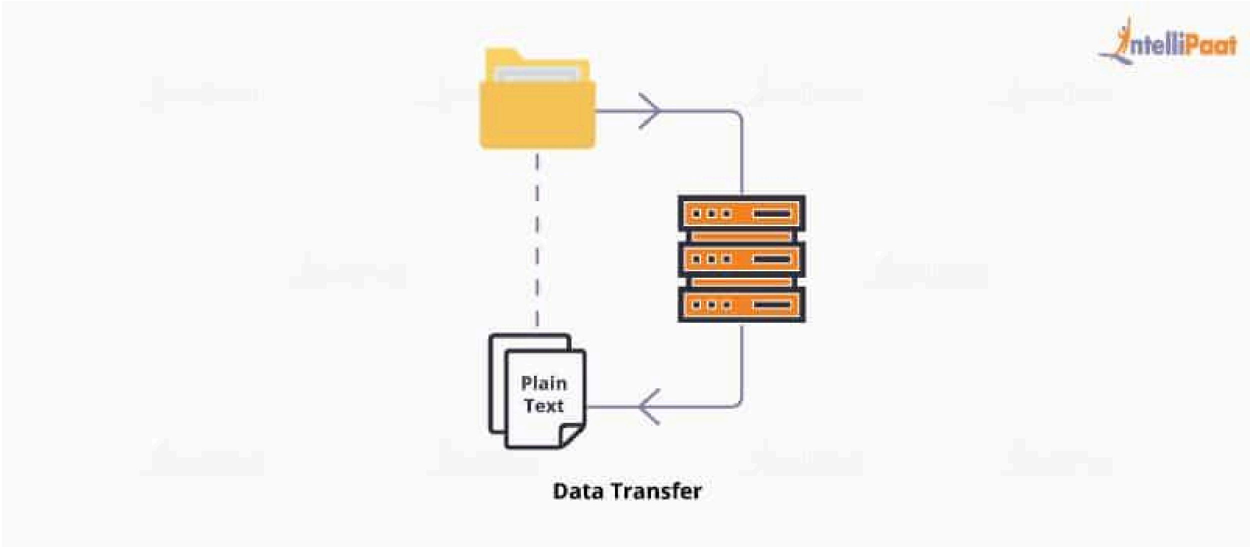

It includes the passing of the login credentials in plain text, i.e., anyone running a sniffer on the network can find the information he/she needs to take control of the device in a few seconds by eavesdropping on a Telnet login session.

38. Differentiate between Linux and Unix

The differences between Linux and Unix are mentioned in the following table:

Linux	Unix
Offers both paid and free OS	Cost varies with the levels
It is portable	It is non-portable
The installation process of Linux does not involve any hardware components	Hardware components are needed to install Unix
It is developed by a worldwide Linux community.	It is developed by AT&T developers.
It is highly flexible and compatible	It is less flexible and compatible compared to Linux.
It is used in both software and hardware, frameworks, etc.	It is used in servers, workstations, etc.

39. Name the four Configuration Management Tools used in UNIX-like operating systems.

- Ansible
- Chef
- Puppet
- CFEngine

40. Mention the difference between BASH and DOS

The key differences between BASH and DOS are given below:

BASH	DOS
Bourne Again Shell	Disk Operating System
Case Sensitive commands	Non Case sensitive commands
/ represents directory separator	/ represents command argument
represents escape character	represents directory separator
Follows conventional naming	Does not follow conventional naming

41. What is the difference between Cron and Anacron?

There are many differences between Cron and Anacron as given below:

- Minimum granularity with Cron is in minutes, while it is in days with Anacron.

- A Cron job can be scheduled by any normal user, while Anacron can be scheduled only by a superuser (a superuser is a special user account used for system administration. Depending on the operating system, the actual name of this account might be root, administrator, admin, or supervisor).
- Cron expects the system to be up and running, while Anacron doesn't expect this all the time. In the case of Anacron, if a job is scheduled and the system is down at this time, it will execute the job as soon as the system is up and running.
- Cron is ideal for servers, while Anacron is ideal for both desktops and laptops.
- Cron should be used when we want a job to be executed at a particular hour and minute, while Anacron should be used when the job can be executed at any time.

42. What are the different run levels in Linux?

In Linux, run levels represent different operating states of a system, each serving a specific purpose. The commonly used run levels are 0 (halt), 1 (single-user mode), 2 (multi-user mode without networking), 3 (multi-user mode with networking), 4 (reserved for user-defined purposes), 5 (multi-user mode with networking and graphical user interface), and 6 (reboot).

43. How do you manage and control processes in Linux using signals?

Managing and controlling processes in Linux is efficiently done through signals. Signals are software interrupts that convey instructions to processes. The 'kill' command is often employed to send signals. For instance, 'kill -9' forcefully terminates a process. 'SIGTERM' (15) requests a graceful termination, while 'SIGHUP' (1) prompts a process to reload its configuration.

44. What is Docker, and how is it useful for Linux administrators?

Docker is a containerization platform that allows Linux administrators to encapsulate applications and their dependencies into containers. This fosters consistency in deployment across different environments. Docker streamlines resource utilization, enhances scalability, and simplifies application management, making it a valuable tool for Linux administrators.

45.What strategies can improve disk I/O performance in Linux?

There are several ways administrators can optimize disk input/output (I/O) speeds in a Linux environment. Using performant file systems like XFS and Btrfs allows for faster reads/writes. Asynchronous I/O with buffered writing can also boost speeds for some workloads. Adjusting elevator algorithms for disk scheduling is another technique. Finally, leveraging solid state drives (SSDs) and RAID can dramatically improve I/O performance.

46. What is logical volume management (LVM), and what advantages does it offer?

Logical volume management (LVM) in Linux provides an abstract layer above physical disks or partitions. This gives administrators the ability to dynamically allocate storage space from a pool of storage resources. Key advantages of LVM include easier resizing of logical volumes, simplified storage provisioning, efficient disk snapshots for backups, and flexible data migration as storage needs change. Overall, LVM introduces welcome storage flexibility for Linux systems.

47. How do you analyze system logs in Linux? What tools can you use?

Analyzing system logs in Linux involves using tools like 'journalctl,' which displays logs from the journal. 'dmesg' shows kernel-related messages, 'tail' allows real-time log monitoring, and 'grep' aids in searching logs for specific patterns. Centralized logging solutions like 'syslog-ng' or 'rsyslog' can aggregate logs for easier analysis.

48. What protocols can provide transport-layer encryption for Linux network traffic?

When securing communications between Linux systems on a network, administrators often employ cryptographic protocols to encrypt data in transit. Options like Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL), secure traffic by creating encrypted tunnels between endpoints. This prevents potential eavesdroppers from accessing transmitted information, essentially guaranteeing transport-level security. Both TLS and SSL rely on underlying public-key infrastructure (PKI) and certificates to facilitate the encrypted sessions. For production Linux environments, staying up-to-date on TLS best practices is advised to maximize network security.

49. What are the best practices for user password security policies in Linux?

Best practices for user password security policies in Linux include enforcing complex passwords, regular password updates, account lockouts after failed attempts, and utilizing tools like 'PAM' (Pluggable Authentication Modules). Implementing multi-factor authentication and restricting root access further enhances security.

50. How can DevOps tools help in automating the deployment of Linux servers?

In the practice of DevOps, there is a lot of emphasis on infrastructure automation, especially when it comes to streamlining the process of deploying Linux servers. For instance, Ansible, Puppet, and Chef are some tools that enable system administrators to program installations with server configurations that are managed through code. The last approach allows for scaling out while keeping uniformity throughout server provisioning. Even better, developers maintain versioning in the infrastructure code, which makes it easier to modify or update systems effectively. Quick deployment times, reduced configuration drift among environments, and infrastructure changes consistent with application development are its advantages over manual server setup. As such, DevOps automation tools make the process less painful by using known patterns.

51. Which high-availability solutions can be used with Linux?

Cluster technologies like Pacemaker and Corosync are some of the existing high-availability solutions in Linux that help to guarantee constant service delivery. Traffic distribution among many servers is enhanced through load-balancing systems such as HAProxy, thus ensuring no single point for failure. Tools like Keepalived also enhance IP failover, hence increasing reliability.

Linux Interview Questions for 4, 5 Years and More Experience

52. What is the name and path of the main system log?

By default, the main system log is /var/log/messages. This file contains all messages and scripts written by a user. By default, all scripts are saved in this file. This is the standard system log file, which contains messages from all system software, non-kernel boot issues, and messages that go to dmesg. The dmesg file is a system file that is written upon the system boot.

Want to be Linux Certified? Learn about Linux Certification!

53. Can we convert a Linux computer into a router in order to enable multiple machines to work on the same Internet connection? If yes, how?

Yes! We can convert a Linux PC into a router so that it can act as an IP gateway for a network. This process of turning a Linux machine into a router is referred to as IP Masquerade, which is basically a Linux networking function that is quite similar to the one-to-many network address translation servers.

Linux IP Masquerading enables the other 'internal' computers that are linked to this Linux system to get connected to the Internet. This Linux feature is available even when these machines do not have their own IP addresses.

In Linux, we can perform IP Masquerading by following the below steps:

Step 1: First of all, we have to make sure that our Linux PC is having an Internet connection, along with a LAN connection. In fact, a Linux PC will be having a PPP connection and an Ethernet card.
Step 2: As the default gateway for TCP/IP networking, all the other systems on our LAN use the Linux machine. Hence, we have to use the same DNS addresses provided by the Internet service provider on all our systems.
Step 3: Now, for enabling IP forwarding, we will use the following command:

```
echo 1 &gt; /proc/sys/net/ipv4/ip_forward
```

For checking whether we have IP forwarding enabled already, we can use the following:

```
sysctl net.ipv4.ip_forward
```

```
net.ipv4.ip_forward = 0
```

Or, we will just check out the value i /proc/sys/net/ipv4/ip_forward:

```
~]# cat /proc/sys/net/ipv4/ip_forward

0
```

Step 4: Finally, we will run /sbin/iptables for setting up those rules that enable IP Masquerading.

Become a Full Stack Web Developer
In collaboration with IBM Microsoft
LEARN MORE
mongoDB.
Express JS
ANGULAR
node

54. What are the different types of modes in VI editor?

The VI editor (Visual Editor) is a basic text editor that appears in most Linux distributions. The following are the main varieties of modes usable in the VI editor:

Command Mode/Regular Mode: The default mode for vi editors is Command Mode/Regular Mode. It is typically used to view and write instructions that perform special or unique vi tasks.

Insertion Mode/Edit Mode: You may use this Insertion mode to edit text or insert text into a file. You can also delete the text.

Ex Mode/Replacement Mode: Ex mode is commonly used for file saving and command execution. We can overwrite the text in this mode.

55. In Linux, how would you change the window manager?

The /.xinitrc file lets us change the window manager that we will use while logging into the X window session. The dot (.) here tells us that this particular file is a hidden. It also means that this file will not be present when we carry out a normal directory listing. In order to set up a window manager, the following command should be saved in this file:

```
exec window manager
```

Once we are done with this part, the next step is to save the file. This way, a new window manager opens up every time we run a startx, and it becomes the default.

Now, let’s check out the commands used for starting some very common desktop environments and window managers:

- KDE = startkde
- GNOME = gnome-session
- BlackBox = blackbox
- FVWM = fvwm
- Window Maker = wmaker
- IceWM = icewm

56. Mention various Linux directory commands

There are five fundamental Linux directory commands for working with files and directories, as listed below:

- pwd: pwd refers to “print working directory”. We use this command to display the path of the current working directory. The syntax of this command is $ pwd.

- cd: cd refers to “change directory”. We use this command to change the present working directory to the specifically required directory. The syntax of this command is $ cd <path to new directory>.
- ls: ls refers to “list”. We use this command to view the full list of files and directories in the present working directory. The syntax of this command is $ ls.
- mkdir: mkdir refers to “make directory”. We use this command to create directories in Linux. The syntax of this command is $ mkdir <name (and path if required) of new directory>.
- rmdir: rmdir refers to “remove directory”. We use this command to remove or delete any directory on the command line. The syntax of this command is $ rmdir <name (and path if required) of directory>.

57. How are shadow passwords given in Linux?

In Linux, the pwconv command is used for providing the shadow passwords. Shadow passwords are given for better system security. This command creates the file /etc/shadow and changes all passwords to ‘x’ in the /etc/passwd file.

First, the entries in the shadowed file, which don’t exist in the main file, are removed. Then, the shadowed entries that don’t have ‘x’ as the password in the main file are updated. Any missing shadowed entries are also added. Finally, passwords in the main file are replaced with ‘x’. These programs can be used for initial conversion as well to update the shadowed file if the main file is edited by hand.

Get certified from the top Linux Course in Singapore now!

58. List out some Linux distributors?

We have a lot of Linux distributors, so we’ll go through a few of the more significant ones.

- Linux Mint: It is a stable and reliable operating system. Mate and Cinnamon are two of the most popular desktop environments used in Linux Mint.
- Debian: It is a Linux distribution that stands for stability, reliability, and a well-oiled release process.
- Manjaro: It provides a pleasurable experience for both novice and seasoned users.
- Ubuntu: Ubuntu is based on Debian and is available in desktop and server variants.
- openSUSE: It is a fantastic choice for both novice and experienced users.

59. What daemon is used for scheduling commands?

The crontab command is used for scheduling commands to run at a later time.

Syntax:

```
crontab [ -u user ] file

crontab [ -u user ] { -l | -r | -e }
```

Options:

- -l: Displays the current crontab entries
- -r: Removes the current crontab
- -e: Edits the current crontab using the editor specified by the VISUAL or EDITOR environment variables

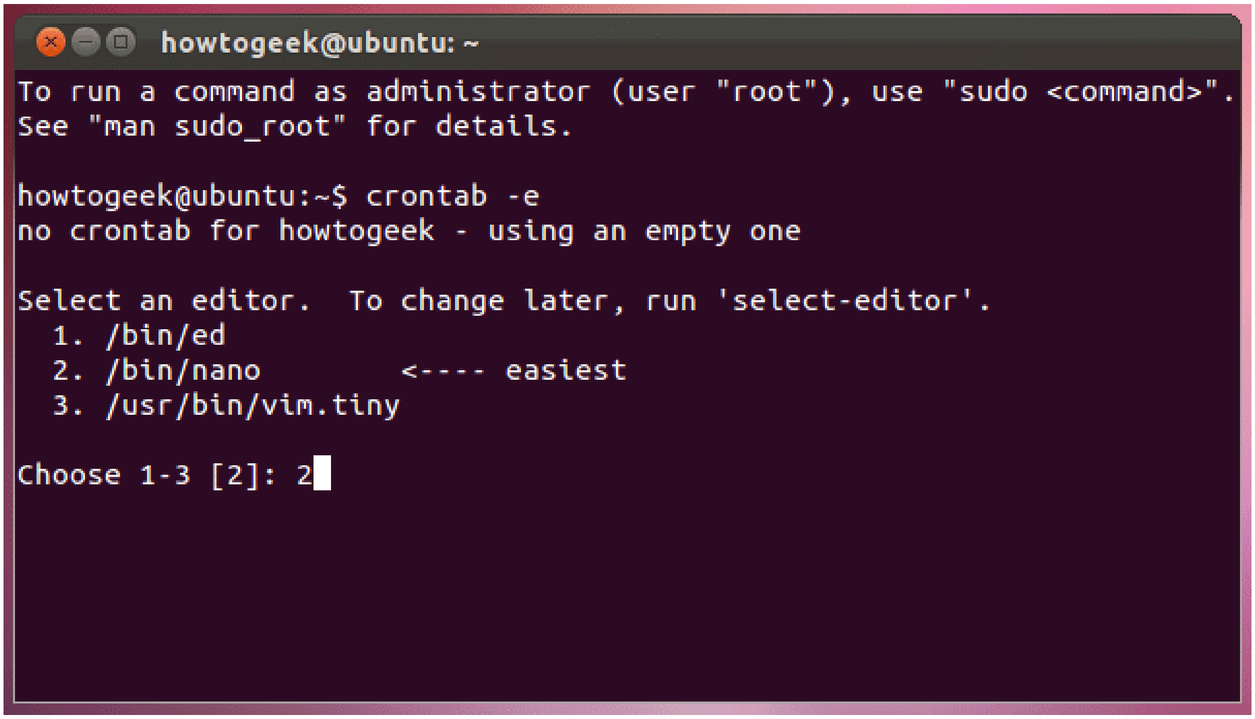

```
To run a command as administrator (user "root"), use "sudo <command>".
See "man sudo_root" for details.

howtogeek@ubuntu:~$ crontab -e
no crontab for howtogeek - using an empty one

Select an editor.  To change later, run 'select-editor'.
  1. /bin/ed
  2. /bin/nano        <---- easiest
  3. /usr/bin/vim.tiny

Choose 1-3 [2]: 2
```

When a user exits from the editor, the modified crontab will be installed automatically. Each user can have their own crontab and though these are files in /var, they are not intended to be edited directly.

If the -u option is given, then the crontab gives the name of the user whose crontab is to be tweaked. If it is given without this option, then it will display the crontab of the user who is executing the command.

60. What do you know about Linux Shell and its types?

The Linux shell is software that serves as a user interface between the user and the kernel. By writing programs, instructions, and scripts on the shell, users may execute instructions and communicate with the kernel. The Linux shell is software that allows users to run commands. It takes human-readable commands as input and translates them into kernel-friendly language.

The main five types of shells in Linux are:

- CSH (C Shell)

- BASH (Bourne Again Shell)
- KSH (Korn Shell)
- TCSH (Tenex C Shell)
- ZSH(Z Shell)

61. What shell does a Linux Administrator assign to a POP3 mail-only account?

A Linux Administrator assigns a POP3 mail-only account to the /bin/false shell. However, assigning a bash shell to a POP3 mail-only account gives the user the login access, which is usually avoided. The /bin/nologin shell can also be used. This shell is provided to the user when we don't want to give shell access to the user. The user cannot access the shell, and it rejects shell login on the server as in Telnet. It is mainly meant for the security of the shells.

POP3 is basically used for downloading mail-to-mail programs. So for the illegal downloading of emails on the shell, this account is assigned to the /bin/false shell or the /bin/nologin shell. Both shells are the same as they do the same work of rejecting the user login to the shell.

The main difference between these two shells is that the false shell shows the incorrect code and any unusual coding when a user logs in to it, whereas the nologin shell simply tells us that no such account is available. Therefore, the nologin shell is used often in Linux.

62. If a volume group named VG0 already exists and we need to extend this volume group up to 4 GB, how do we do it?

To extend an existing volume group named VG0 to a size of 4 GB, you would typically follow these steps:

1. Check the available space: Use the vgdisplay command to check the available space in the VG0 volume group. Look for the “Free PE / Size” field to determine how much free space is currently available.

```
vgdisplay VG0
```

2. If there is not enough free space in the VG0 volume group, you may need to add physical volumes (disks) to increase the available space. This step assumes that you have additional disks available. If you already have sufficient free space, you can skip to step 3.

- Add physical volumes: Use the pvcreate command to initialize the additional disks and make them available for use in the volume group.

```
pvcreate /dev/sdb1      # Replace /dev/sdb1 with the
appropriate device name for your disk
```

- Extend volume group: Use the vgextend command to add the newly created physical volumes to the VG0 volume group.

```
vgextend VG0 /dev/sdb1  # Replace /dev/sdb1 with the
appropriate device name for your disk
```

3. Extend the logical volume: Once you have sufficient free space in the VG0 volume group, you can extend the logical volume (LV) within it. Use the lvextend command to increase the size of the LV to 4 GB.

```
lvextend -L 4G /dev/VG0/LVname   # Replace LVname with the
name of the logical volume you want to extend
```

4. Resize the file system: Finally, you need to resize the file system on the LV to utilize the newly allocated space. The specific command depends on the file system type.

```
For ext2/ext3/ext4 file systems, use the resize2fs command.

resize2fs /dev/VG0/LVname   # Replace LVname with the name of the logical volume

For XFS file systems, use the xfs_growfs command.

xfs_growfs /dev/VG0/LVname   # Replace LVname with the name of the logical volume
```

After completing these steps, your VG0 volume group should be extended to 4 GB, and the file system on the logical volume should be resized to utilize the additional space.

63. Is there any relation between the modprobe.conf file and network devices?

Yes, this file assigns a kernel module to each network device.

Example:

```
[root@localhost ~]# cat /etc/modprobe.conf

alias eth0 b44
```

```
Here, b44 is the kernel module for network device eth0.

We can confirm whether this module “b44” is present or not by
the following command

[root@localhost ~]# lsmod |grep b44

b44 29005 0
```

64. How are hardening and security policies implemented in Linux?

There are multiple processes involved in implementing security measures and hardening Linux. Start by using package managers like “yum” or “apt” to update the system with security fixes regularly. Turn off unused services, filter network traffic with firewalls, and set up AppArmor or SELinux for the required access constraints. Limit user access, enforce strict password requirements, and keep an eye out for questionable activity in system logs.

65. How can the performance of a Linux web server be optimized?

Choosing effective web server software, such as Apache or Nginx, is one way to optimize web server performance on Linux. Use caching systems to lessen server demand, such as Redis or Varnish. Use Content Delivery Networks (CDNs), optimize code, and configure server settings. Putting SSL/TLS into practice for safe communication and load

66. What are the different types of firewall architectures used in Linux networks?

Linux networks commonly utilize two main types of firewall architectures: Stateful and Stateless. Stateful firewalls track the state of active connections, allowing only legitimate traffic based on the connection's state. Stateless firewalls filter packets based on predefined rules without considering the connection state. Tools like iptables and firewalld are often used to implement these architectures.

67. How does Linux implement Access Control Lists (ACLs) for permissions?

Linux implements Access Control Lists (ACLs) to grant or deny permissions beyond the traditional owner, group, and other settings. ACLs allow specifying permissions for specific users or groups on a file or directory. The 'setfacl' and 'getfacl' commands are used to manage and display ACLs, providing finer-grained control over access permissions.

68. What tools enable container orchestration and management on Linux platforms?

Running containerized workloads at scale on Linux often requires orchestration software to automate deployment, networking, scaling, and availability. Popular open-source options include Kubernetes, which has become a standard for production-grade container orchestration. Other tools like Docker Swarm and Red Hat OpenShift also help streamline managing containers across clusters of Linux hosts. These leverage API-driven configurations to roll out microservices-based applications packaged in Docker or OCI compliant runtimes. Best practices involve infrastructure definitions within YAML or JSON files for the desired state. Overall, Linux container orchestration hinges on automation for operational efficiency across on-premise or cloud environments.

69. How can Linux systems connect to Active Directory for centralized user authentication?

Enabling Active Directory authentication for Linux servers involves cross-platform single sign-on software. Common open source options are SSSD (System Security Services Daemon) or Winbind. These run on the Linux side, allowing the hosts to join Windows Server domains. Once configured, Linux can redirect user login attempts against AD's Kerberos-based authentication and directory services. This ties into Linux PAM (Pluggable Authentication Modules) to enable domain users seamless access to Linux resources using the same credentials. Overall, mature open source tools bridge the interoperability gap for unified access control between Active Directory and Linux systems.

70. What are the best practices for storage management in Linux?

Best practices for storage management in Linux include using Logical Volume Management (LVM) for dynamic volume control, monitoring disk space regularly, and employing filesystem quotas. Implementing backup and recovery strategies, optimizing I/O operations, and selecting appropriate RAID configurations contribute to effective storage management.

71. How can you optimize Linux for SSD drives?

Optimizing Linux for SSD drives involves using filesystems with native SSD support, such as ext4 or XFS. Align partitions properly, enable TRIM support for SSDs, and adjust mount options for better performance. Utilize the 'fstrim' command to maintain SSD health and consider tweaking I/O scheduler settings for optimal SSD performance.

72. What tools can help sysadmins configure and manage Linux servers?

Handling server configurations at scale poses challenges. Thankfully, Linux administrators can leverage configuration management software to codify and automate management. Widely adopted open source options include Ansible, Puppet, and Chef. These define desired infrastructure states using declarative definitions, rather than scripts. Once deployed, they can install packages, tweak system files, start services, and otherwise configure Linux boxes based on version controlled recipes. This enhances consistency while reducing server drift across environments. Configuration management is a pillar of the DevOps movement, benefiting stability, security, and operational efficiency within dynamic Linux environments.

73.How can Linux server monitoring integrate with public cloud platforms?

Monitoring Linux server performance often involves sending telemetry to cloud provider platforms for consolidated dashboards and alerts. On AWS, CloudWatch Agents extract OS and application metrics from EC2 instances. Similarly, Google Cloud Platform stacks run monitoring agents for workload visibility. Platforms like Azure Monitor pull Linux stats to drive auto-scaling rules. Open source tools like Prometheus also tie cloud-based Linux metrics into higher level observability pipelines. Overall, modern monitoring relies on agents and federated data to correlate metrics across dynamic cloud and container deployments.

Linux Commands Interview Questions

74. What is YUM?

YUM stands for Yellowdog Updater Modified because it is based on YUP, the Yellowdog Updater. Yellowdog is a version of Linux for the Power Architecture hardware and is RPM-based, just like Red Hat Enterprise Linux and Fedora. YUP, and later YUM, were written by the Linux community as a way to maintain an RPM-based system.

75. What is the role of Kudzu?

Kudzu is used to detect new hardware. Red Hat Linux runs a hardware discoverer, called Kudzu. When attempting to identify a serial port, Kudzu resets the serial port. This stops the serial console. Kudzu is configured from the following file:

```
/etc/sysconfig/kudzu
```

Kudzu can be prevented from resetting hardware, by setting the configuration parameter SAFE to 'yes.'

76. What is the difference between ext2 and ext3 file systems?

- The ext3 file system is an enhanced version of the ext2 file system.
- The most important difference between ext2 and ext3 is that ext3 supports journaling.
- After an unexpected power failure or system crash (also called an unclean system shutdown), each ext2 file system must be checked for consistency by the e2fsck program. This is a time-consuming process and during this time, any data on the volumes is unreachable.
- The journaling provided by the ext3 file system means that this sort of a file system check is no longer necessary after an unclean system shutdown. The only time a consistency check occurs while using ext3 is in certain rare hardware failure cases, such as hard drive failures. The time to recover an ext3 file system after an unclean system shutdown does not depend on the size of the file system or on the number of files. Rather, it depends on the size of the journal used to maintain consistency. The default journal size takes almost a second to recover, depending on the speed of the hardware.

77. Explain the /proc file system?

The /proc file system is a virtual file system that provides detailed information about Linux Kernel, hardware, and running processes. Files under the /proc directory are named as virtual files.

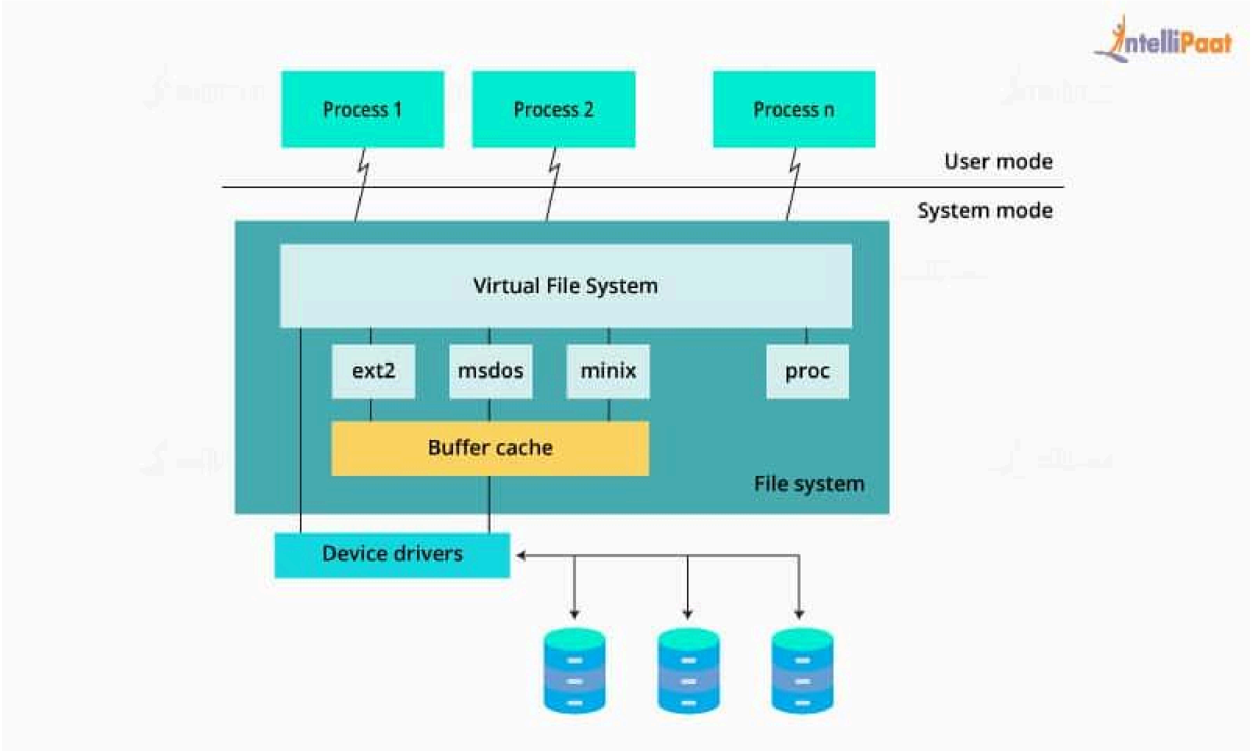

Since /proc contains virtual files, it is called a virtual file system. These virtual files have unique qualities. Most of them are listed as zero bytes in size. Virtual files such as /proc/interrupts, /proc/meminfo, /proc/mounts, and /proc/partitions provide an up-to-the-moment glimpse of the system's hardware. Others, such as the /proc/filesystems file and the /proc/sys/ directory, provide system configuration information and interfaces.

78. In Linux, how can I figure out where a file is stored?

To find the path to the file, use the locate command. If you wish to locate the locations of a file named sample.txt, use the following command:

```
$ locate sample.txt
```

79. How would you create an ext4 file system?

We can create an ext4 file system with the following command:

```
# mke2fs -t ext4 /dev/DEV
```

80. In Linux, how do you stop a running process?

Every process has its own identifier. We must first locate the process id in order to terminate it. The “ps” command displays a list of all currently active processes, along with their ids. The “kill” command is then used to end the process.

81. How to enable ACLs for the /home partition?

To enable Access Control Lists (ACLs) for the /home partition, you’ll need to follow these general steps:

1. Check if ACLs are already enabled: Use the mount command to check if the /home partition is mounted with the acl option. Look for an entry that includes acl in the options field.

```
mount | grep /home
```

If you see an entry that includes acl in the options field, it means ACLs are already enabled. You can skip to step 4.

2. Backup important data: Before making any changes, it’s always a good practice to back up important data on the /home partition. This step helps ensure that you have a copy of the data in case anything goes wrong during the process.

3. Edit the /etc/fstab file: To enable ACLs permanently, you need to modify the /etc/fstab file.

- Open the /etc/fstab file using a text editor. For example:

```
sudo nano /etc/fstab
```

- Locate the entry for the /home partition. It may look similar to:

```
UUID=&lt;partition-UUID&gt;   /home   &lt;filesystem-type&gt;   defaults   0   2
```

- Modify the options field to include acl. It should look like:

```
UUID=&lt;partition-UUID&gt;   /home   &lt;filesystem-type&gt;   defaults,acl   0   2
```

Make sure to replace <partition-UUID> with the actual UUID of the /home partition and <filesystem-type> with the appropriate file system type.

Save the changes and exit the text editor.

4. Remount the partition: After modifying the /etc/fstab file, you need to remount the /home partition for the changes to take effect.

- Run the following command to remount the partition:

```
sudo mount -o remount /home
```

5. Verify ACLs are enabled: Use the mount command again to verify that the /home partition is now mounted with the acl option.

```
mount | grep /home
```

If you see an entry that includes acl in the options field, it means ACLs are now enabled for the /home partition.

Once ACLs are enabled, you can use commands like getfacl and setfacl to manage access control lists on files and directories within the /home partition.

82. How can Linux administrators concatenate multiple files into a single file?

A common task when managing Linux servers involves combining data from separate files into one unified file. The ‘cat’ utility provides a handy way to link together contents sequentially. For instance, to merge text across two files named file1.txt and file2.txt into a target merged_file.txt:

cat file1.txt file2.txt < merged_file.txt[/code] This concatenates file1, followed by file2, storing the resulting data stream into the merged target. The greater than ‘>’ character effectively redirects merged output. Understanding basic POSIX commands like cat enables administrators to efficiently wrangle Linux files and data.

83. How do you find files containing a particular string pattern recursively in Linux?

To find files containing a particular string pattern recursively in Linux, you can use the ‘grep’ command along with the ‘r’ option for recursive search. For example,

```
grep -r "pattern" /path/to/search
```

This command recursively searches for the specified “pattern” in files under the specified directory (/path/to/search) and displays the matching lines along with the file names.

84. How can Linux administrators run commands at a later time?

Linux systems provide the ‘at’ utility for administrators to schedule tasks in the future without cron’s repetition. For example, to set a job to execute a single time at 2:30 PM tomorrow:

```
at 2:30pm tomorrow
```

After confirming the time, ‘at’ will prompt for the actual command, script, or pipeline to run at that point. This allows delayed execution without dedicated cron entries or sleep commands. Helpful scenarios include planning updates during low-traffic hours, staggering resource-intensive tasks, or automating notifications based on events. Overall, the at command gives administrators valuable control over one-off executions.

85. How can Linux admins run commands as other users without switching sessions?

Linux administrators often need to briefly “become” other users to perform tasks requiring different permissions. Instead of fully logging in and out, the ‘sudo’ command offers a convenient way to run specific commands as alternate users on-the-fly. For example, to quickly edit a file owned by the user ‘johndoe’:

```
sudo -u johndoe vim /home/johndoe/documents/letter.txt
```

After entering the current admin password, this would open letter.txt to edit as johndoe without starting a new shell session. Sudo authorization provides flexibility to operate as required users only when necessary. Understanding user switching with sudo is invaluable for Linux permissions management without elevated accounts.

86. How can Linux admins view only subdirectories within the current folder?

Navigating Linux directories frequently involves checking subfolder contents. Instead of manual scanning, the ls command supports filtering directory-only listings. For example, to output only child directories in long format:

```
ls -ld */
```

The -d flag displays directories themselves rather than their contents, while the wildcard symbol */ matches subdirectories. This provides a streamlined way to view nested filesystem structure changes from the current working path. Understanding versatile ls options helps Linux admins efficiently traverse and manage directories at scale.

87. How can Linux administrators analyze text file contents and size?

When managing configurations, log files, or other text documents on Linux servers, admins often need to check details like length and quantity. The ‘wc’ (word count) utility handles these content statistics easily. To print newline, word, and byte counts for a file:

```
wc filename.txt
```

This summarizes size metrics in order, great for confirming expectations or identifying outliers. ‘wc’ can also accept input pipes and multiple files at once for convenient aggregation. Understanding basic wc transforms abstract text data into actionable admins insights. Whether checking web logs, user records, or API payloads, wc provides a simple method to inspect important Linux file metadata.

88. How can Linux admins filter directory listings to avoid specific matches?

When browsing Linux directories, admins may want to avoid clutter from certain subfolders like logs, caches, or temp files. Rather than manual inspection, the ls command can filter results by piping output into grep. For example, to hide dirs like tmp/ and log/:

```
ls -l | grep -vE 'tmp|log'
```

The -v flag inverts grep's matches, while -E enables extended expressions. This displays a long list format while skipping specified directories. Chaining POSIX file commands like ls, grep, awk and find lets admins elegantly parse and transform directory content. Understanding these patterns unlocks more selective navigation for Linux shell scripting.

89. How do you empty or delete the contents of a file without removing the file in Linux?

To empty or delete the contents of a file without removing the file in Linux, you can use the 'truncate' command with the '–size' option. For example:

```
truncate --size 0 filename
```

This command sets the size of the specified file ('filename') to 0, effectively emptying its contents.

90. How can Linux admins segment large files into smaller chunks?

Applications often produce massive log or data files that create storage and transfer challenges. Rather than clumsy workarounds, the split command reliably breaks files into manageable pieces. For example, to divide largefile.txt into 1 megabyte parts:

```
split -b 1M largefile.txt split_file
```

This sequentially segments the data, using split_file as the prefix for numbered chunks. Split handles file system limitations gracefully while accommodating pipelines. Understanding file slicing tools like split helps Linux admins wrangle and control outsized file outputs.

91. How can Linux admins find symlink metadata and destinations?

Managing Linux servers often involves reviewing symbolic link (symlink) purpose and origins. While ls displays arrows denoting symlinks, the target destinations are hidden. Instead, find reveals details with:

```
find /path -type l -ls
```

This recursively searches for symlink files, reporting inode information and linkage targets using -ls. Intelligent find parameters help admins expose symlink ownership, permissions, and aim points for maintenance or auditing needs.

Linux Admin Interview Questions

92. How can administrators monitor and optimize Linux disk I/O?

Tracking disk I/O helps troubleshoot Linux performance issues. Tools like iotop identify busy processes in real-time, while iostat graphs activity over time. Blktrace captures request-level traces for analysis. Tuning the I/O scheduler algorithm via elevator kernel parameters can optimize request queuing. Swapping to faster storage like SSDs dramatically accelerates read/write speeds as well. Robust monitoring combined with a high performance disk subsystem results in responsive Linux infrastructure.

93. What buying considerations influence Linux server hardware decisions?

Selecting optimized hardware for Linux deployments depends on workloads. Evaluation includes processor core count for computational needs, RAM capacity for memory-intensive apps, and storage volume for data retention. Network interface bandwidth and connectivity enable fast client access. Virtualization support, scalability options, and energy efficiency also factor in. Considering both current Linux resource demands and future growth, it is imperative for ideal procurement.

94. What tools and strategies help admins manage Linux capacity?

Monitoring overall system health via sar provides historical usage trends for right-sizing needs. Granular memory stats from vmstat assist in tuning applications and planning growth. Checking disk utilization with df flags storage shortage risks. Combining predictive monitoring, scaling up resources, moving workloads to cloud, and archiving old data are key capacity strategies.

95. How can Infrastructure as Code automate Linux cloud provisioning?

Tools like Terraform and CloudFormation codify the Linux stack for cloud, enabling declarative deployment automation. Admins define desired infrastructure states using config files or templates for rapid, reproducible environments. Platforms like AWS further support scripted security groups and network builds. IaC replaces manual clicking with programmatic Linux instance provisioning.

96. What considerations facilitate robust Linux disaster recovery?

Guarding against Linux service disruption relies on full-system backups using rsync or object storage for offsite redundancy. Orchestration tools can script failover to standby instances. Testing plan execution validates recovery time objectives. Combining frequent backups with accessible, warm spare capacity ensures Linux availability despite outages.

97. How can Linux integrate with developer source code tools?

Linking Linux infrastructure to coding involves version control systems like Git. Repositories on the Linux host mirror code branches, while SSH keys secure developer access. Git hooks trigger automation as authors commit changes. CI platforms like Jenkins then compile and deploy application artifacts in a streamlined SDLC pipeline. Deep code integration accelerates delivery on Linux.

98. What factors influence container vs. VM adoption on Linux?

Deciding between Linux containers and VMs depends on resource constraints, speed, security, and app characteristics. Containers provide lightweight resource isolation using the host kernel, ideal for microservices. Rapid provisioning suits CI/CD pipelines. However, VMs offer stronger separation for multi-tenant and legacy apps, through higher overhead. Evaluating workload requirements and infrastructure strategies guides the optimal Linux virtualization approach.

99. How do admins evaluate in-house vs. cloud Linux hosting tradeoffs?

Key criteria for Linux hosting decisions include control, cost, and staffing models. In-house infrastructure grants oversight but demands upfront CapEx and ongoing management. Cloud's OpEx pricing brings flexibility and automation but risks vendor lock-in. Compliance and data sovereignty considerations also influence strategies. Organizations mix hybrid solutions to maximize Linux hosting advantages while mitigating downsides.

100. What patch management best practices help secure Linux systems?

Keeping Linux systems updated with security patches involves automating package managers like yum or apt. Scheduling regular scans and installs ensures fix availability. Maintaining a mirrored test environment validates update stability before hitting production. Employing tools like Unattended Upgrades for hands-off hardening and documenting patch policy protects against threats while avoiding disruption.

101. How can admins diagnose bottlenecks in Linux database performance?

Pinpointing sluggish Linux database servers begins with system monitoring using top, htop, and iostat for CPU, memory, and disk resources. Database tooling like mysqladmin further isolates activity by tables, queries, and connections. Analyzing slow query logs, execution plans and tuning indexes optimizes data access. Holistically correlating Linux and database telemetry exposes performance gaps, guiding optimization for smooth database operations.

Linux Troubleshooting Interview Questions

102. What steps can you take if your Linux server is having trouble resolving domain names?

If your Linux server is facing DNS issues, begin by examining the DNS server configuration in '/etc/resolv.conf.' Use commands like 'nslookup' or 'dig' to query DNS servers, check network connectivity, and review firewall settings. Additionally, investigate DNS logs in '/var/log' for errors and employ 'systemctl' to troubleshoot or restart the DNS resolver service.

103. How do you investigate and address the startup failure of an Nginx web server on a Linux system after a reboot?

When an Nginx web server on Linux fails to start post-reboot, troubleshooting involves checking error logs in '/var/log/nginx,' validating the Nginx configuration with 'nginx -t,' and inspecting system logs using 'journalctl.' Assess resource availability, including disk space and memory. If necessary, manually start Nginx with verbose debugging ('nginx -g "daemon off;"') to pinpoint specific issues preventing startup.

104. What steps can be taken to resolve

To resolve "Disk quota exceeded" errors for a Linux user, utilize the 'quota' command to review the user's quota limits. Adjust quotas using 'edquota' if needed, delete unnecessary files, and identify large files or directories using 'du' and 'find.' If the issue persists, ensure proper disk space allocation and consider expanding quotas or disk space.

105. How can you identify the root cause of occasional freezes on your Linux server, requiring a hard reboot?

For occasional Linux server freezes, analyze system logs (‘dmesg,’ ‘syslog,’ ‘journalctl’) for error messages and patterns leading up to the freeze. Monitor system resources with ‘top’ or ‘htop’ to detect potential resource exhaustion. Check for kernel panics, hardware issues, or problematic applications. Tools like ‘sar’ for historical system activity can also aid in diagnosis

106. What measures would you take to recover from a Linux kernel panic after a kernel update and prevent its recurrence?

To recover from a kernel panic post-update, reboot the system and choose an older, stable kernel from the bootloader menu. For prevention, ensure hardware and installed modules are compatible before updating the kernel. Maintain a backup of critical data, regularly update the system, and stay informed about potential issues related to specific kernel updates through community forums.

107. How do you troubleshoot situations where an application encounters permission errors preventing it from opening files in Linux?

Troubleshooting permission errors in Linux when an application can’t open files involves checking file permissions with ‘ls -l’ and verifying ownership. Confirm that the user running the application has the necessary permissions. Use ‘strace’ or ‘ltrace’ to trace system calls and library calls made by the application, gaining insights into permission-related issues. Adjust file permissions or ownership as needed.

108. What steps would you take if a scheduled Linux backup job encountered broken pipe errors?

In the event of broken pipe errors in a scheduled Linux backup job, the initial course of action involves scrutinizing the backup script, examining backup logs for errors, and confirming the availability and disk space of the backup destination. Ensure the accessibility of the receiving end (e.g., remote server) and troubleshoot any network issues. Thoroughly review the script’s error-handling mechanisms and address potential intermittent connectivity problems.

109. How do you troubleshoot a scenario where a Linux machine is unable to connect to specific websites?

Troubleshooting a Linux machine that can’t connect to specific websites requires a systematic approach. Begin by checking network connectivity using ‘ping’ or ‘traceroute,’ examine DNS resolution with ‘nslookup’ or ‘dig,’ and review firewall settings. Rule out the impact of proxy configurations on internet access. Test connections to other websites to isolate the issue and investigate potential DNS or routing problems, making necessary adjustments to network settings.

110. How would you debug a custom-written Linux daemon process that occasionally crashes?

Debugging a custom-written Linux daemon process prone to occasional crashes involves a multifaceted approach. Utilize debugging tools like ‘gdb’ to analyze core dumps or attach to the running process. Enable debugging symbols during compilation for more informative stack traces. Scrutinize logs for error messages and timestamps, monitor system resources, and delve into the daemon’s source code to identify potential bugs. Implement additional logging for specific events to facilitate the debugging process.

111. What steps do you take to troubleshoot a Linux machine that occasionally freezes during boot?

Troubleshooting a Linux machine experiencing occasional freezes during boot necessitates a thorough investigation. Check boot logs in '/var/log' for errors, leverage 'journalctl' to analyze system logs, and inspect the bootloader configuration, such as GRUB. Disable unnecessary services during boot using 'systemctl,' identify problematic startup processes, and verify hardware integrity. Examine for disk issues and apply available updates. Utilize single-user mode to initiate a minimal boot environment for isolating and resolving the issue.

Linux Networking Interview Questions

112. How do you activate packet forwarding in Linux to route network traffic between interfaces?

To enable packet forwarding in Linux for routing network traffic between interfaces, follow these steps:

1. Edit the sysctl configuration file:

```
sudo nano /etc/sysctl.conf
```

2. Uncomment or add the line:

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

3. Apply the changes:

```
sudo sysctl -p
```

4. Temporarily enable forwarding:

```
sudo sysctl net.ipv4.ip_forward=1
```

113. Your Linux router encounters issues connecting to specific public IP addresses. How would you troubleshoot this?

To troubleshoot a Linux router that is unable to connect to specific public IP addresses, consider the following:

1. Verify router connectivity using ‘ping’:

```
ping [public IP addresses]
```

2. Inspect the routing table:

```
ip route show
```

3. Review firewall rules:

```
iptables -L
```

4. Check DNS resolution:

```
nslookup [or] dig
```

5. Investigate relevant logs in ‘/var/log/syslog’ or ‘journalctl’.

114. How do you configure Linux as a DHCP server for dynamically assigning IP addresses to machines on a LAN?

- Install DHCP server: `sudo apt-get install isc-dhcp-server` (for Ubuntu/Debian)
- Configure DHCP settings in `/etc/dhcp/dhcpd.conf`
- Specify network interface in `/etc/default/isc-dhcp-server`
- Restart DHCP server: `sudo service isc-dhcp-server restart`

115. What Linux tools are available for network traffic analysis?

Linux users have access to several tools for analyzing network traffic:

A sophisticated graphical packet analyzer is **Wireshark:**.

A command-line packet sniffer is **tcpdump:**.

– **load:** gives an easy-to-read overview of all incoming and outgoing communications.

– **iftop:** is a console-based network bandwidth meter operating in real-time.

– **iperf:** evaluates the performance of TCP and UDP.

116. How can Linux be set up to distribute requests across several web servers in the manner of a software load balancer?

To configure Linux as a load balancer for software:

1. Set up a load balancing program such as Nginx or HAProxy.

2. Adjust the load balancer's parameters, including the balancing algorithms and backend servers.
3. Review the setup thoroughly and keep an eye on the load balancer logs.

117. You need to apply QoS policies to prioritize business-critical traffic on your Linux network. How will you implement this?

- Use `tc` (traffic control) command to configure QoS policies
- Set up classes and filters to prioritize traffic based on criteria
- Define bandwidth limits and priorities for different classes
- Monitor QoS statistics using `tc` and adjust settings as needed

118. How do you configure Linux as a VPN server for establishing secure remote access tunnels?

To configure Linux as a VPN server,

1. Choose OpenVPN or IPsec for VPN implementation.
2. Install and configure the chosen VPN server software.
3. Generate certificates and configure user authentication.
4. Set up firewall rules to allow VPN traffic.
5. Start and enable the VPN server service.

119. How can Linux be integrated with an intrusion detection system (IDS) for network monitoring?

Integrating Linux with an IDS involves:

1. Install and configure an IDS (e.g., Snort or Suricata).
2. Define rules for detecting suspicious network activity.
3. Integrate with syslog for centralized logging.

4. Regularly update IDS rules for the latest threat intelligence.
5. Analyze IDS alerts and take appropriate action.

120. What steps are involved in aggregating multiple network interfaces into a single logical interface in Linux?

To aggregate multiple network interfaces in Linux:

1. Load the bonding kernel module: `sudo modprobe bonding`.
2. Configure bonding in '/etc/network/interfaces' or an equivalent file.
3. Specify the bonding mode (e.g., active-backup, balance-rr).
4. Activate the bonded interface: `sudo ifup bond0`.

121. How do you capture and analyze packet traces using tcpdump for troubleshooting networking issues in Linux?

To capture and analyze packet traces with tcpdump:

1. Install tcpdump: `sudo apt-get install tcpdump`.
2. Capture packets: `sudo tcpdump -i <interface> -w <output_file.pcap>`.
3. Analyze with Wireshark or tcpdump filters.
4. Use specific options for protocol-specific analysis.
5. Understand source/destination IP, ports, and protocol details.

Linux Professionals' Salary Trends

In the US, the salary range of a Linux administrator is between $82,000 and $1,00,000, with an average of $1,05,441 per year. In India, the salary range of a

Linux engineer is between ₹3,00,000 and ₹10,00,000 per year, with an average of ₹5,40,000 per year.

Job Role	Salary Range
Linux Administrator	₹4L – ₹7L/yr
Linux Engineer	₹3L – ₹10L/yr

Linux Professionals’ Job Trends

- Global Demand: There is increasing demand for Linux administrators and engineers to manage Linux-based infrastructure and cloud environments. Companies like AWS and Azure make extensive use of Linux. There are more than 97,000 job postings for Linux in the United States on LinkedIn.
- Projected Growth: According to the BLS, between 2022 and 2032, IT jobs will rise by 15%, adding 46,900 jobs annually.
- Regional Trends: Rise in Linux jobs for DevOps engineers with Linux skills to enable CI/CD pipelines and containerization using Docker and Kubernetes. Linux forms the base for many modern cloud-native stacks. There are more than 20,313 jobs in India posted for Linux on Naukri.com.

Linux Professionals’ Roles & Responsibilities

As posted in an actual job description by Oracle, we can summarize a few common responsibilities for a linux professional.

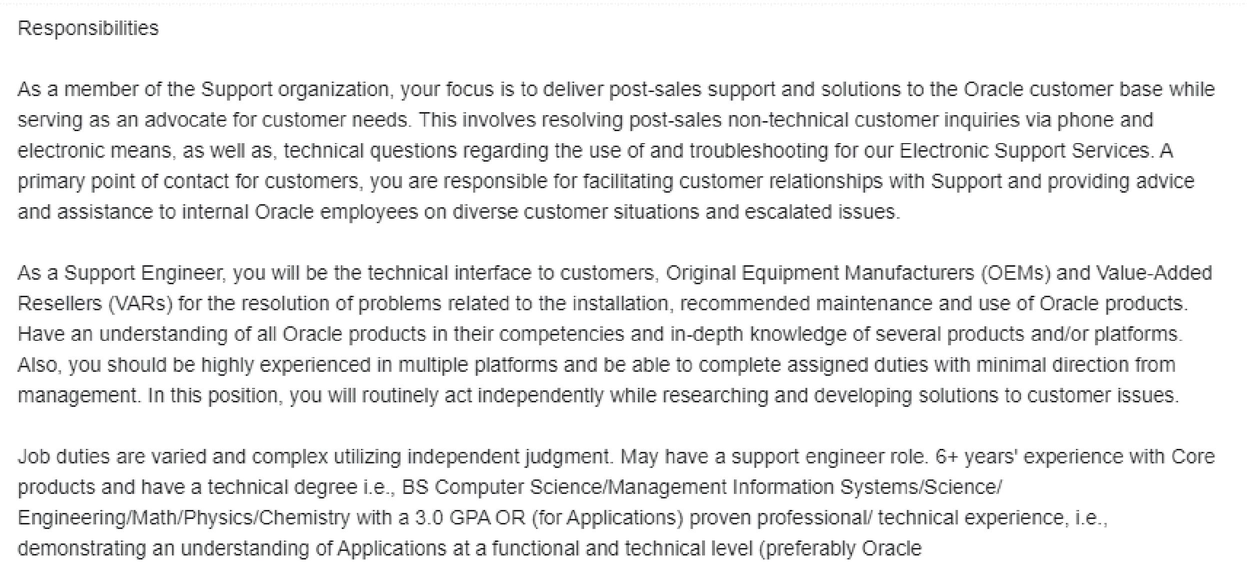

Responsibilities

As a member of the Support organization, your focus is to deliver post-sales support and solutions to the Oracle customer base while serving as an advocate for customer needs. This involves resolving post-sales non-technical customer inquiries via phone and electronic means, as well as, technical questions regarding the use of and troubleshooting for our Electronic Support Services. A primary point of contact for customers, you are responsible for facilitating customer relationships with Support and providing advice and assistance to internal Oracle employees on diverse customer situations and escalated issues.

As a Support Engineer, you will be the technical interface to customers, Original Equipment Manufacturers (OEMs) and Value-Added Resellers (VARs) for the resolution of problems related to the installation, recommended maintenance and use of Oracle products. Have an understanding of all Oracle products in their competencies and in-depth knowledge of several products and/or platforms. Also, you should be highly experienced in multiple platforms and be able to complete assigned duties with minimal direction from management. In this position, you will routinely act independently while researching and developing solutions to customer issues.

Job duties are varied and complex utilizing independent judgment. May have a support engineer role. 6+ years' experience with Core products and have a technical degree i.e., BS Computer Science/Management Information Systems/Science/ Engineering/Math/Physics/Chemistry with a 3.0 GPA OR (for Applications) proven professional/ technical experience, i.e., demonstrating an understanding of Applications at a functional and technical level (preferably Oracle

The most common responsibilities of a professional in Linux are as follows:

1. Scripting and automation for various tasks.
2. Security and integrity of the systems.
3. Troubleshooting problems and finding a way to optimally resolve the issue.
4. Backup the important information using the various tools at hand.
5. Performance and monitoring of each and every process to optimize the entire workflow.

Conclusion

We hope this set of Linux interview questions will help you prepare for your interviews. We wish you luck!

Enroll today in our comprehensive Linux Certification Training to start your career or enhance your skills in the field of Linux and get certified today.

If you're eager to explore additional Linux interview questions in depth, feel free to join Intellipaat's Linux Community and get answers to your queries.