

Unix/Linux Basic interview questions

Question 1: What is the major advantage of a hash table? (Asked by Silicon Magic Corp. people)

Answer: The major advantage of a hash table is its speed. Because the hash function is to take a range of key values and transform them into index values in such a way that the key values are distributed randomly across all the indices of a hash table.

Question 2: What are the techniques that you use to handle the collisions in hash tables?(Asked by Silicon Magic Corp. people)

Answer: We can use two major techniques to handle the collisions. They are open addressing and separate chaining. In open addressing, data items that hash to a full array cell are placed in another cell in the array. In separate chaining, each array element consist of a linked list. All data items hashing to a given array index are inserted in that list.

Question 3: In Unix OS, what is the file server? (Asked by Silicon Magic Corp. people)

Answer: The file server is a machine that shares its disk storage and files with other machines on the network.

Question 4: What is NFS? What is its job?(Asked by Silicon Magic Corp. people)

Answer: NFS stands for Network File System. NFS enables filesystems physically residing on one computer system to be used by other computers in the network, appearing to users on the remote host as just another local disk.

Question 5: What is CVS? List some useful CVS commands.(Asked by Silicon Magic Corp.people)

Anser: CVS is Concurrent Version System. It is the front end to the RCS revision control system which extends the notion of revision control from a collection of files in a single directory to a hierarchical collection of directories consisting of revision controlled files. These directories and files can be combined together to form a software release. There are some useful commands that are being used very often. They are

cvsexec checkout
cvsexec update
cvsexec add
cvsexec remove
cvsexec commit

Unix/Linux administration interview questions

What is LILO?

LILO stands for Linux boot loader. It will load the MBR, master boot record, into the memory, and tell the system which partition and hard drive to boot from.

What is the main advantage of creating links to a file instead of copies of the file?

A: The main advantage is not really that it saves disk space (though it does that too) but, rather, that a change of permissions on the file is applied to all the link access points. The link will show permissions of lrwxrwxrwx but that is for the link itself and not the access to the file to which the link points. Thus if you want to change the permissions for a command, such as su, you only have to do it on the original. With copies you have to find all of the copies and change permission on each of the copies.

Write a command to find all of the files which have been accessed within the last 30 days.

```
find / -type f -atime -30 > December.files
```

This command will find all the files under root, which is '/', with file type is file. '-atime -30' will give all the files accessed less than 30 days ago. And the output will put into a file call December.files.

What is the most graceful way to get to run level single user mode?

A: The most graceful way is to use the command `init s`.
If you want to shut everything down before going to single user mode then do `init 0` first and from the ok prompt do a `boot -s`.

What does the following command line produce? Explain each aspect of this line.

```
$ (date ; ps -ef | awk '{print $1}' | sort | uniq | wc -l ) >> Activity.log
```

A: First let's dissect the line: The `date` gives the date and time as the first command of the line, this is followed by the a list of all running processes in long form with UIDs listed first, this is the `ps -ef`. These are fed into the `awk` which filters out all but the UIDs; these UIDs are piped into `sort` for no discernible reason and then onto `uniq` (now we see the reason for the `sort` - `uniq` only works on sorted data - if the list is A, B, A, then A, B, A will be the output of `uniq`, but if it's A, A, B then A, B is the output) which produces only one copy of each UID.

These UIDs are fed into `wc -l` which counts the lines - in this case the number of distinct UIDs running processes on the system. Finally the results of these two commands, the date and the `wc -l`, are appended to the file "Activity.log". Now to answer the question as

to what this command line produces. This writes the date and time into the file Activity.log together with the number of distinct users who have processes running on the system at that time. If the file already exists, then these items are appended to the file, otherwise the file is created.

How would you make the following SQL statement run faster? SELECT * FROM TABLEA WHERE COL1='A' AND COL2='B';

Make sure that COL1 and COL2 have indexes.
Find out which condition will return less values and use that as the first conditional.

What is Data Mining

Data mining is the process of sifting through extremely large amounts of Data to find trends or relevant information.

Name the Seven layers in the OSI Model.

Application, Presentation, Session, Transport, Network, Data Link, Physical

What is one way to view a unix network share on a Windows computer, within explorer

NFS, The Unix computer can be running a NFS Server Daemon.

How would you find all the processes running on your computer.

Unix, is ps -ef or ps -aux depending on version.

What is DHCP

DHCP is a way to dynamically assign IP address to computers. Dynamic Host Configuration Protocol

What is HTTP Tunneling

HTTP Tunneling is a security method that encrypts packets traveling through the internet. Only the intended recipient should be able to decrypt the packets. Can be used to Create Virtual Private Networks. (VPN)

You have 9 identical looking balls, however one ball is heavier than the others. You have two chances to use a balance. How do you find out which ball is the heaviest?

Split into groups of three, randomly choose two groups and use balance on them. If one group is heavier, then discard the other 6 balls. If the two groups are the same weight. The heavier ball must be in the group that was not on the scale. Now randomly choose

two balls and test on balance. If they are the same weight, the heaviest ball is on one that was not tested. Else the heaviest ball is already known from the balance.