System Calls Job Interview Questions And Answers

About Interview Questions Answers

Interview Questions Answers . ORG is an interview preparation guide of thousands of Job Interview Questions And Answers, Job Interviews are always stressful even for job seekers who have gone on countless interviews. The best way to reduce the stress is to be prepared for your job interview. Take the time to review the standard interview questions you will most likely be asked. These interview questions and answers on System Calls will help you strengthen your technical skills, prepare for the interviews and quickly revise the concepts.

If you find any **question or answer** is incorrect or incomplete then you can **submit your question or answer** directly with out any registration or login at our website. You just need to visit System Calls Interview Questions And Answers to add your answer click on the *Submit Your Answer* links on the website; with each question to post your answer, if you want to ask any question then you will have a link *Submit Your Question*; that's will add your question in System Calls category. To ensure quality, each submission is checked by our team, before it becomes live. This System Calls Interview preparation PDF was generated at **Wednesday 29th November, 2023**

You can follow us on FaceBook for latest Jobs, Updates and other interviews material. www.facebook.com/InterviewQuestionsAnswers.Org

Follow us on Twitter for latest Jobs and interview preparation guides. https://twitter.com/InterviewQA

If you need any further assistance or have queries regarding this document or its material or any of other inquiry, please do not hesitate to contact us.

Best Of Luck.

Interview Questions Answers.ORG Team
https://InterviewQuestionsAnswers.ORG/
Support@InterviewQuestionsAnswers.ORG

System Calls Interview Questions And Answers Guide.

Question - 1:

What happens when we do insmod & rmmod in Linux Device Drivers?

Ans:

insmod: insmod is a tool used to attach a module to the
running linux kernel. This will take the kernel object(.ko)
and takes all executable code and data sections of the .ko
and attach it to the running linux kernel.
rmmod: used to remove or deattach a module code from the
running kernel

View All Answers

Question - 2:

What kind of information the Linux driver modules (.ko) files has?

Ans:

kernel 2.6 introduces a new file naming convention: kernel
modules now have a .ko extension (in place of the old .o
extension) which easily distinguishes them from conventional
object files. The reason for this is that they contain an
additional .modinfo section that where additional
information about the module is kept.
Linux program modpost can be used to convert .o files into
.ko files.

View All Answers

Question - 3:

What are the different ways the Linux can switch from User Space to Kernel Space & vice-versa?

Ans:

There are 2 situations when Linux can switch from user Space
to Kernel Space:-
1) by doing System calls
2) When interrupt comes (to handle interrupt)
3) by executing 128 (0x80) instruction or doing sysenter
Linux can switch from kernel Space to User space:-
1) process in kernel mode is preempted.
2) After completion of Interrupt handler / System call
3) performing sysexit sys call

View All Answers

Question - 4:

What is stored in /proc?

Ans:

Mainly hardware related information such as CPU
information, Memory (RAM) information stored under /proc
directory
example:
cat /proc/cpuinfo (show the information of CPU of that
particular hardware)
cat /proc/meminfo (show the information of Memory i.e.
RAM of that particular hardware)

View All Answers

Question - 5:

What is stored at /lib/modules?

Ans:

It contains all the kernel modules that needed to be loaded
into kernel (booting etc). there will some .map, .dep
(dependency files) files present.
When the kernel needs a feature that is not resident in the
kernel, the kernel module daemon kmod[1] execs modprobe to
load the module in.
You can see what modules are already loaded into the kernel
by running lsmod, which gets its information by reading the
file /proc/modules

View All Answers

Question - 6:

What is atomic function and atomic variable?

Ans:

atomic variables are the variables which can only be
manipulated atomically using atomic APIs. Linux declares
variable as atomic by using the type atomic_t. Basically
used a way to achieve synchronization.
an atomic operation is one which cannot be (or is not)
interrupted by concurrent operations and cannot be broken up
into smaller parts that could be performed by different
processors.
Atomic function is a function which is executed to
completion without interruption. Atomic function can also be
seen as a small critical section which is executed without
interruption, locking.

View All Answers

Question - 7:

Linux file defaults permition is?

Ans:

umask value = 022
Without a umask in effect,any file created will have 666
permissions.
  666
  022

  644

A umask of 022 will result in files created with 666 permission.

View All Answers

Question - 8:

How to create secured appeche web sever?

Ans:

You need to install an SSL certificate in apahce to secure
the transactions.

View All Answers

Question - 9:

What do fork() internally call?

Ans:

Linux implements fork() via the clone() system call.
The clone() system call, in turn, calls do_fork().
The bulk of the work in forking is handled by do_fork(),
which is defined in kernel/fork.c.This function calls
copy_process() and then starts the process running.
If copy_process() returns successfully, the new child is
woken up and run. Deliberately, the kernel runs the child
process first.

View All Answers

Question - 10:

What does exec family return?

Ans:

When successful exec will not return, it will start
executing the new program

However if there is an- error exec returns -1 and sets the
errno to the appropriate value

View All Answers

Question - 11:

What is difference between spinlock, seamaphores and mutex and where to use it?

Ans:

mainly spinlock used in threads to avoid synchronization,where as semaphore and mutex used to avoid process synchronization.
1.spinlock is something like polling.it spins for resouce until aloocated resouce releases.
2.binary semaphore and mutex are similar.

View All Answers

Question - 12:

How to find out the dependency required for a package?

Ans:

#rpm -qpR filename.rpm
Lists the dependency list of packages.

View All Answers

Question - 13:

What is the diff between ssh and telnet?

Ans:

ssh is secured shell, allows the user to login remotely with more secured.
whereas telnet also same but authentications like passwords, transfers over a network as text mode. so it is not good to use.

View All Answers

Question - 14:

How to create swap partition after OS installation?

Ans:

swap can be created in two ways after the installation,
1. fdisk command
2. create a swap file using dd command
after creating swap file or file system
#mkswap /dev/sda10
#swapon /dev/sda10
#swapon -s   #To see the swap devices
by using dd command
#dd if=/dev/zero of=/swap bs=1024 count=1
Which will creates the file size 1024(1GB).
#mkswap /swap
#swapon /swap
#swapon -s   #to see the swap devices

View All Answers

Question - 15:

If the FS is in read-only mode, so we cannot create any file. How will you fix it?

Ans:

LVM is a mechanism use for providing specality of extending (or) reducing
the sizes of an existing partition.

View All Answers

Question - 16:

How to use resize2fs, what is the purpose?

Ans:

resize2fs is only for ext2 filesystem but not ext3.
 first unmount the partition
 #umount /dev/sda1

 #tune2fs -O ^has_journal /dev/sda1   #to remove journal from /dev/sda1

 #e2fsck -f /dev/sda1

 #resize2fs /dev/sda1 600M   #resize the partition

View All Answers

Question - 17:

What is the diff between ext3 and ext2 File system?

Ans:

ext3 is also same as the ext2, but journaling concept is introduced in ext3.
Compared to ext2, ext3 is slow. ext2 less secure compared to ext3. ext2 is less
Performance where as ext3 is very good performance.
View All Answers

Question - 18:

How to remove a PV from lvm without any data loss?

Ans:

by using pvremove command.
View All Answers

Question - 19:

What is the largest disk size can be used in LVM?

Ans:

Don't know exactly, think of 2TB or 8TB
View All Answers

Question - 20:

Difference between Raid 1 and Raid 5?

Ans:

RAID 1 is disk striping. no mirroring no parity. Minimum 2 disks required. If any
One disk fails all the data get lost.
   RAID 5 is disk striping with parity. Minimum 3 disks required. if anyone disk fails
Data is safe, if two fails data get lost.
View All Answers

Question - 21:

How will you harden the server?

Ans:

A Server-- it is weather in testing or production-- are primary targets for
the attackers. By taking the proper steps, you can turn a vulnerable box into
a hardened server.
 How to secure SSH sessions, configure firewall rules, minimize software, listed below,
 1. Encrypt Data communication
   -- use scp, ssh avoid FTP, Telnet and Rlogin /rsh
 2. Minimize Software to minimize vulnerability
   -- use RPM pkg management / YUM utility to remove unwanted packages installed
 3. One Network Service per System or Vm Instance
   -- Run different network services on separate servers or vm instance.
 For example, if an attacker able to successfully exploit software called
 Apache flow, he/she get an access to entire server including other services
 such as MYSQL, email server and so on.
 4. Keep linux software and Kernel up to date.
   -- Use yum update or up2date
 some distros apt-get update
 5. Security essentials like selinux
 6. password authentication like password aging, restricting to user previous
 passphrases, and locking user accounts after login failures.
 7. Disable unwanted services using chkconfig --list | grep "3:on"
View All Answers

Question - 22:

Where the History file can be located?

Ans:

.bash_history
View All Answers

Question - 23:

How to configure sendmail server on red hat Linux version 4 and what and all we Require?

Ans:

Sendmail should be installed by default when you install Red
Hat Linux. If it is not then you need to install the
Sendmail RPM's with the Red Hat distribution.
Configuring Sendmail
----
1) Edit file "/etc/mail/sendmail.mc" - Look for the line:
   DAEMON_OPTIONS(`Port=smtp,Addr=127.0.0.1, Name=MTA')

Change this line to:
dnl DAEMON_OPTIONS(`Port=smtp,Addr=127.0.0.1, Name=MTA')
Save the file.
2)Make the sendmail configuration file:
m4 /etc/mail/sendmail.mc > /etc/sendmail.cf
3)Restart Sendmail:
/etc/rc.d/init.d/sendmail restart

View All Answers

Question - 24:

What is nis server?

Ans:

NIS is a service that provides any user on that network
with the same working environment irrespective of the
system on that network which has been used for login
purpose.
For example if NIS server is set up in a single system and
configured to hold user accounts and their passwords and
access information. Then any user on that network can login
to his/her account that is set up in the NIS server from
any system (with nis client running) on that configured
network. This gives a look and feel that the user is logged
into his/her own system. But actually its the account on
the NIS server that is mounted on the local sytem on user
login .

View All Answers

Question - 25:

What is samba, what is configuration file, how it will work?

Ans:

Samba provides file and print services to all manner of
SMB/CIFS clients, including the numerous versions of
Microsoft Windows operating systems. Samba configuration
file is smb.conf:
Sample smb.conf

[global]
    # Domain name ..
    workgroup = DOMAIN.NAME
    # Server name - as seen by Windows PCs ..
    netbios name = SERVER1
    # Be a PDC ..
    domain logons = Yes
    domain master = Yes
    # Be a WINS server ..
    wins support = true
    # allow user privileges
    #enable privileges = yes
    obey pam restrictions = Yes
    dns proxy = No
    os level = 35
    log file = /var/log/samba/log.%m
    max log size = 1000
    syslog = 0
    panic action = /usr/share/samba/panic-action %d
    pam password change = Yes
    # Allows users on WinXP PCs to change their password
when they press Ctrl-Alt-Del
    unix password sync = no
    ldap passwd sync = yes
    # Printing from PCs will go via CUPS ..
    load printers = yes
    printing = cups
    printcap name = cups
    # Use LDAP for Samba user accounts and groups ..
    passdb backend = ldapsam:ldap://localhost
    # This must match init.ldif ..
    ldap suffix = dc=domain,dc=name
    # The password for cn=admin MUST be stored in
/etc/samba/secrets.tdb
    # This is done by running 'sudo smbpasswd -w'.
    ldap admin dn = cn=admin,dc=domain,dc=name
    # 4 OUs that Samba uses when creating user accounts,
computer accounts, etc.
    # (Because we are using smbldap-tools, call them
'Users', 'Computers', etc.)
    ldap machine suffix = ou=Computers
    ldap user suffix = ou=Users

```
        ldap idmap suffix = ou=Idmap
        # Samba and LDAP server are on the same server in
this example.
        ldap ssl = no
        # Scripts for Samba to use if it creates users,
groups, etc.
        add user script = /usr/sbin/smbldap-useradd -m '%u'
        delete user script = /usr/sbin/smbldap-userdel %u
        add group script = /usr/sbin/smbldap-groupadd -p '%g'
        delete group script = /usr/sbin/smbldap-groupdel '%g'
        add user to group script =
/usr/sbin/smbldap-groupmod -m '%u' '%g'
        delete user from group script =
/usr/sbin/smbldap-groupmod -x '%u' '%g'
        set primary group script = /usr/sbin/smbldap-usermod
-g '%g' '%u'
        # Script that Samba users when a PC joins the domain ..
        # (when changing 'Computer Properties' on the PC)
        #add machine script = /usr/sbin/smbldap-useradd -w '%u'
        add machine script = /usr/sbin/useradd -s /bin/false
-d /home/nobody %u
        # Values used when a new user is created ..
        # (Note: '%L' does not work properly with
smbldap-tools 0.9.4-1)
        logon drive = H:
        logon home = \server%U
        logon path = \serverProfiles%U
        logon script = logon.bat
        # This is required for Windows XP client ..
        server signing = auto
        server schannel = Auto
[homes]
        comment = Home Directories
        path = /home/users/%U
        valid users = %S
        read only = No
        browseable = No
[netlogon]
        comment = Network Logon Service
        path = /var/lib/samba/netlogon
        admin users = root
        guest ok = Yes
        browseable = No
[Profiles]
        comment = Roaming Profile Share
        # would probably change this to elsewhere in a
production system ..
        path = /var/lib/samba/profiles
        read only = No
        profile acls = Yes
        browsable = No
        hide files = /desktop.ini/ntuser.ini/NTUSER.*/
[printers]
        comment = All Printers
        path = /var/spool/samba
        use client driver = Yes
        create mask = 0600
        guest ok = Yes
        printable = Yes
        browseable = No
        public = yes
        writable = yes
        admin users = root
        write list = root
[print$]
        comment = Printer Drivers Share
        path = /var/lib/samba/printers
        write list = root
        create mask = 0664
        directory mask = 0775
        admin users = root
Test it with :
# testparm /etc/samba/smb.conf
Load smb config files from /etc/samba/smb.conf
Processing section "[homes]"
Processing section "[netlogon]"
Processing section "[Profiles]"
Processing section "[printers]"
Processing section "[print$]"
Loaded services file OK.
Server role: ROLE_DOMAIN_PDC
Press enter to see a dump of your service definitions
```

Question - 26:

What is iptabe on RedHat Linux?

Ans:

Iptables is the userspace command line program used to configure the Linux 2.4.x and 2.6.x IPv4 packet filtering ruleset. Iptables allows administrators to configure the operating system so that it allows applications and clients to connect through the network and stop unwanted applications and clients from communicating and corrupting the operating system.
It is not specific to Redhat. It is available in all linux 2.4.x and 2.6.x kernels.

Question - 27:

What is sudo on Linux?

Ans:

The sudo command stands for "superuser do". If a server needs to be administered by a number of people it is normally not a good idea for them all to use the root account. This is because it becomes difficult to determine exactly who did what, when and where if everyone logs in with the same credentials. The sudo utility was designed to overcome this difficulty.
The sudo utility allows users defined in the /etc/sudoers configuration file to have temporary access to run commands they would not normally be able to due to file permission restrictions. The commands can be run as user "root" or as any other user defined in the /etc/sudoers configuration file.
The privileged command you want to run must first begin with the word sudo followed by the command's regular syntax. When running the command with the sudo prefix, you will be prompted for your regular password before it is executed.
You may run other privileged commands using sudo within a five-minute period without being re-prompted for a password.
All commands run as sudo are logged in the log file /var/log/messages.
In order to use sudo we first need to configure the sudoers file.
Do not edit directly the file:
To edit it, use the command
visudo
******Output***************
/etc/sudoers
#
This file MUST be edited with the 'visudo' command as root.
#
See the man page for details on how to write a sudoers file.
#
Defaults   env_reset
Host alias specification
User alias specification
Cmnd alias specification
User privilege specification
root   ALL=(ALL) ALL
****
You will see the line
root ALL=(ALL) ALL
This lines means that the user root can execute from ALL terminals, acting as ALL (any) users, and run ALL (any) command.
The first part is the user, the second is the terminal from where the user can use sudo, the third is as which user he may act, and the last one, is which commands he may run.
Example:
Granting Access To Specific Users To Specific Files
----
amsin21, %operator ALL= /sbin/, /usr/sbin,
/usr/local/apps/check.pl
This entry allows user amsin21 and all the members of the group operator to gain access to all the program files in the /sbin and /usr/sbin directories, plus the privilege of running the command /usr/local/apps/check.pl. Notice how the trailing slash (/) is required to specify a directory location:
Granting Access to Specific Files as Another User

The sudo -u entry allows allows you to execute a command as

if you were another user, but first you have to be granted
this privilege in the sudoers file.
This feature can be convenient for programmers who sometimes
need to kill processes related to projects they are working
on. For example, programmer amsin21 is on the team
developing a financial package that runs a program called
monthend as user accounts. From time to time the application
fails, requiring "amsin21" to stop it with the /bin/kill,
/usr/bin/kill or /usr/bin/pkill commands but only as user
"accounts". The sudoers entry would look like this:
amsin21 ALL=(accounts) /bin/kill, /usr/bin/kill, /usr/bin/pkill
User amsin21 is allowed to stop the monthend process with
this command:
sudo -u accounts pkill monthend
Granting Access Without Needing Passwords

This example allows all users in the group operator to
execute all the commands in the /sbin directory without the
need for entering a password. This has the added advantage
of being more convenient to the user:
%operator ALL= NOPASSWD: /sbin/
Using Aliases in the sudoers File

Sometimes you'll need to assign random groupings of users
from various departments very similar sets of privileges.
The sudoers file allows users to be grouped according to
function with the group and then being assigned a nickname
or alias which is used throughout the rest of the file.
Groupings of commands can also be assigned aliases too.
In the next example, users amsin21, amsin211 and amsin212
and all the users in the operator group are made part of the
user alias ADMINS. All the command shell programs are then
assigned to the command alias SHELLS. Users ADMINS are then
denied the option of running any SHELLS commands and su:
Cmnd_Alias   SHELLS = /usr/bin/sh, /usr/bin/csh,
              /usr/bin/ksh, /usr/local/bin/tcsh,
              /usr/bin/rsh, /usr/local/bin/zsh

User_Alias   ADMINS = amsin21, amsin211, amsin212, %operator
ADMINS   ALL   = !/usr/bin/su, !SHELLS
This attempts to ensure that users don't permanently su to
become root, or enter command shells that bypass sudo's
command logging. It doesn't prevent them from copying the
files to other locations to be run. The advantage of this is
that it helps to create an audit trail, but the restrictions
can be enforced only as part of the company's overall
security policy.

View All Answers

Question - 28:

How to configure http server on red hat linux4.0?

Ans:

Simple configuration

Install apache: (get the rpm for apache)
rpm -ivh httpd.rpm
vi /etc/httpd/conf/httpd.conf
and change
ServerName your-domain.com
Start apache : service httpd start
Verify by pointing the browser to http://localhost/

View All Answers

Question - 29:

What are the backup utilities on red hat Linux 4.0?

Ans:

cpio with find command ,tar with gun/bunzip ,dump with only
ext2 & ext3 Fs used to take backups
Also you have amanda software in RHEL 4.0 version to take a
backup

View All Answers

Question - 30:

What is cups and how to configure?

Ans:

CUPS (formerly an acronym for Common Unix Printing System, but now with no official expansion) is a modular printing system for Unix-like computer operating systems which allows a computer to act as a print server. A computer running CUPS is a host that can accept print jobs from client computers, process them, and send them to the appropriate printer.
If you are using a client with CUPS and a CUPS server has already been configured, installing the printers on your client can not get much easier than this: do nothing. Through broadcasting, the client should find the CUPS server and automatically configure the printers that are installed on that print server. This is one of the features of CUPS that will be really appreciated on large networks.
Manually configuring printers with CUPS, also is a peace of cake. If you are new to CUPS and/or Unix printing, the way to go is probably the web interface. If you have to configure lots of printers, using the command-line will probably be faster.
The URL to access the CUPS web interface is http://hostname:631/admin by default. The port can be changed in cupsd.conf if necessary.
To add a printer from the command-line the general syntax is lpadmin -p printer -E -v device -m ppd Lpadmin with the -p option adds or modifies a printer. The printers are saved in the file The -x option deletes the named printer. Read the lpadmin man page for available options.
Example 3. command-line examples
/usr/sbin/lpadmin -p testpr1 -E -v socket://192.168.1.9 -m deskjet.ppd
/usr/sbin/lpadmin -p testpr2 -E -v parallel:/dev/lp0 -m laserjet.ppd
/usr/sbin/lpadmin -x testpr1
More information about configuring printers and options can be found in the CUPS documentation. The Software Administrators Manual will teach you all you need to know about configuring printers with CUPS.

View All Answers

Question - 31:
How to create a ftp user on RedHat Linux 4.0?

Ans:

firstly create the user.
useradd surendra
passwd surendra
then opne the FTP file:-
vi /etc/vsftpd/vsftpd.conf
add this line at the end of the file
userlist_deny=NO
then open this file,
vi /etc/vsftpd/user_list
and add the above created user in this file to access FTP services.
and restart the FTP services
/etc/init.d/vsftpd start
chkconfig vsftpd on (for permanent on)

View All Answers

Question - 32:
How to create samba server in fedora Linux 9?

Ans:

vi /etc/samba/smb.conf
Add these lines
[fedora 9]
comment=publicstuff
path=/share
public=yes
writable=yes
available=yes
create mask=0777
directory mask=0777

View All Answers

Question - 33:
Why the kernel panic error was appearing?

Ans:

(1)Suppose any problem in initrd file then kernel panic error showing

(2)Any file system problem or any change or remove file system then showing the kernel panic error

View All Answers

Question - 34:

what is mean by raid and what are all raids available even in software and hardware?

Ans:

Raid is Redundant Array of Independent Disks/Device.It is
Technology to improve Disk read & write Performence and
FaultTolerance., By adding new disk u can recover data if
one of disk goes down / fails. parity is a calculated
technique to rebuild data from disk fails.
levels in Raid
0 - stripping
1 - Mirroring
3 - Striping with Parity
5 - Striping with Parity with more fault tolerant widley used.

View All Answers

Question - 35:

How to get microseconds of system time from Redhat Linux 4.0?

Ans:

The systemcall "gettimeofday" can be used to get the time in microseconds. The call takes two arguments.
1- struct timeval
2- struct timezone (for timezone information). you can have second argument null.
timeval structure has two fields:
tv_sec (represents time in seconds)
tv_usec (represents time in microseconds)
the code snapshot is :
////
struct timeval t_time;
gettimeofday(&t_time,NULL);

View All Answers

Linux OS Most Popular & Related Interview Guides

1 : **Device Drivers Interview Questions and Answers.**

2 : **Linux OS Management Interview Questions and Answers.**

3 : **Linux Makefile Interview Questions and Answers.**

4 : **Linux Environment Interview Questions and Answers.**

5 : **Linux OS Shell Interview Questions and Answers.**

6 : **GCC Compiler Interview Questions and Answers.**

7 : **Linux OS Interview Questions and Answers.**

8 : **Bash Arithmetic Expressions Interview Questions and Answers.**

9 : **Linux IPC Interview Questions and Answers.**

10 : **Linux General Interview Questions and Answers.**