

RedHat Interview Questions Adminsitrator

Q: - How are devices represented in UNIX?

All devices are represented by files called special files that are located in /dev directory.

Q: - What is 'inode'?

All UNIX files have its description stored in a structure called 'inode'. The inode contains info about the file-size, its location, time of last access, time of last modification, permission and so on. Directories are also represented as files and have an associated inode.

Q: - What are the process states in Unix?

As a process executes it changes state according to its circumstances. Unix processes have the following states:

Running : The process is either running or it is ready to run .

Waiting : The process is waiting for an event or for a resource.

Stopped : The process has been stopped, usually by receiving a signal.

Zombie : The process is dead but have not been removed from the process table.

Q: - What command should you use to check the number of files and disk space used and each user's defined quotas?

repquota

Q: - What command is used to remove the password assigned to a group?

gpasswd -r

Q: - What can you type at a command line to determine which shell you are using?

echo \$SHELL

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

find / -type f -atime -30 > filename.txt

Q: - What is a zombie?

Zombie is a process state when the child dies before the parent process. In this case the structural information of the process is still in the process table.

Q: - What daemon is responsible for tracking events on your system?

syslogd

Q: - What do you mean a File System?

File System is a method to store and organize files and directories on disk. A file system can have different formats called file system types. These formats determine how the information is stored as files and directories.

Q: - Tell me the name of directory structure hierarchy for Linux

/root

/boot

/bin

/sbin

/proc

/mnt

/usr
/var
/lib
/etc
/dev
/opt
/srv
/tmp
/media

Q: - What does /boot directory contains?

The /boot/ directory contains static files required to boot the system, such as the Linux kernel, boot loader configuration files. These files are essential for the system to boot properly.

Q: - If some one deletes /boot directory from your server, than what will happen?

In that case your server will be in unbootable state. Your Server can't boot without /boot directory because this directory contains all bootable files

Q: - What does /dev directory contain?

The /dev directory contains all device files that are attached to system or virtual device files that are provided by the kernel.

Q: - What is the role of udev daemon?

The udev demon used to create and remove all these device nodes or files in /dev/ directory.

Q: - What kind of files or nodes /dev/ directory contains and how do I access or see device files?

Block Device Files:-

Block device files talks to devices block by block [1 block at a time (1 block = 512 bytes to 32KB)].

Examples: - USB disk, CDROM, Hard Disk

```
# ls /dev/sd*
```

```
brw-rw---- 1 root root 8, 0 Mar 15 2009 sda  
brw-rw---- 1 root root 8, 1 Mar 15 2009 sda1  
brw-rw---- 1 root root 8, 2 Mar 15 2009 sda2  
brw-rw---- 1 root root 8, 3 Mar 15 2009 sda3  
brw-rw---- 1 root root 8, 4 Mar 15 2009 sda4  
brw-rw---- 1 root root 8, 16 Mar 15 2009 sdb
```

Character Device Files:-

Character device files talk to devices character by character.

Examples: - Virtual terminals, terminals, serial modems, random numbers

```
#ls /dev/tty*
```

```
crw-rw---- 1 root root 4, 64 Mar 15 2009 ttyS0  
crw-rw---- 1 root root 4, 65 Mar 15 2009 ttyS1  
crw-rw---- 1 root root 4, 66 Mar 15 2009 ttyS2  
crw-rw---- 1 root root 4, 67 Mar 15 2009 ttyS3
```

Q: - Tell me the name of device file for PS/2 mouse connection.

/dev/psaux

Q: - Tell me the name of device file for parallel port (Printers).

/dev/lp0

Q: - What does /etc/X11/ directory contains?

The /etc/X11/ directory is for X Window System configuration files, such as xorg.conf.

Q: - What does /etc/skel directory contains?

The /etc/skel directory contains files and directories that are automatically copied over to a new user's home directory when such user is created by the useradd or adduser command.

Q: - Tell me name of Linux File systems?

Ext2

Ext3

Q: - 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 mounted ext2 file system on the machine 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 file system check is no longer necessary after an unclean system shutdown. The only time a consistency check occurs 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 the number of files; rather, it depends on the size of the journal used to maintain consistency. The default journal size takes about a second to recover, depending on the speed of the hardware.

Q: - Any idea about ext4 file system?

The ext4 or fourth extended filesystem is a journaling file system developed as the successor to ext3. Ext4 filesystem released as a functionally complete and stable filesystem in Linux with kernel version 2.6.28.

Features of ext4 file system:-

1. Currently, Ext3 supports 16 TB of maximum file system size and 2 TB of maximum file size.

Ext4 have 1 EB of maximum file system size and 16 TB of maximum file size.

[An EB or exabyte is 1018 bytes or 1,048,576 TB]

2. Fast fsck check than ext3

3 In Ext4 the journaling feature can be disabled, which provides a small performance improvement.

4. Online defragmentation.

5. Delayed allocation

Ext4 uses a filesystem performance technique called allocate-on-flush, also known as delayed allocation. It consists of delaying block allocation until the data is going to be written to the disk, unlike some other file systems, which may allocate the necessary blocks before that step.

Q: - How we create ext3 file system on /dev/sda7 disk?

```
# mkfs -j /dev/sda7
```

Q: - Can we convert ext2 filesystem to ext3 file system?

Yes, we can convert ext2 to ext3 file system by tune2fs command.
tune2fs -j /dev/

Q: - How we will create ext4 file system?

mke2fs -t ext4 /dev/DEV

Q: - Explain /proc filesystem?

/proc is a virtual filesystem that provides detailed information about Linux kernel, hardware's and running processes. Files under /proc directory named as Virtual files. Because /proc contains virtual files that's why it is called 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, like the /proc/filesystems file and the /proc/sys/ directory provide system configuration information and interfaces.

Q: - Can we change files parameters placed under /proc directory?

Yes

To change the value of a virtual file, use the echo command and a greater than symbol (>) to redirect the new value to the file. For example, to change the hostname on the fly, type:
echo <http://www.nextstep4it.com> > /proc/sys/kernel/hostname

Q: - What is the use of sysctl command?

The /sbin/sysctl command is used to view, set, and automate kernel settings in the /proc/sys/ directory.

Q: - /proc/ directory contains a number of directories with numerical names. What is that?

These directories are called process directories, as they are named after a program's process ID and contain information specific to that process.

Q: - What is RAID?

RAID, stands for Redundant Array of Inexpensive Disks. RAID is a method by which same data or information is spread across several disks, using techniques such as disk striping (RAID Level 0), disk mirroring (RAID Level 1), and disk striping with parity (RAID Level 5) to achieve redundancy, lower latency, increased bandwidth, and maximized ability to recover from hard disk crashes.

Q: - Why should we use RAID?

System Administrators and others who manage large amounts of data would benefit from using RAID technology.

Following are the reasons to use RAID

- Enhances speed
- Increases storage capacity using a single virtual disk
- Minimizes disk failure

Q: - What is the difference between hardware RAID and Software RAID?

The hardware-based RAID is independent from the host. A Hardware RAID device connects to the SCSI controller and presents the RAID arrays as a single SCSI drive. An external RAID system moves all RAID handling "intelligence" into a controller located in the external disk subsystem. The whole subsystem is connected to the host via a normal SCSI controller and appears to the host as a single disk.

Software RAID is implemented under OS Kernel level. The Linux kernel contains an MD driver

that allows the RAID solution to be completely hardware independent. The performance of a software-based array depends on the server CPU performance and load.

Q: - What are the commonly used RAID types?

RAID 0
RAID 1
RAID 5

Q: - Explain RAID 0?

RAID level 0 works on “striping” technique. In RAID 0 the array is broken down into strips and data is written into strips. RAID 0 allows high I/O performance but provides no redundancy. RAID 0 Array Size is equal to sum of disks in array. If one drive fails then all data in the array is lost.

Q: - Explain RAID 1?

RAID Level 1 is based on Mirroring technique. Level 1 provides redundancy by writing identical data to each member disk of the array. The storage capacity of the level 1 array is equal to the capacity of one of the mirrored hard disks in a Hardware RAID or one of the mirrored partitions in a Software RAID. RAID 1 provides redundancy means good protection against disk failure. In RAID 1 write speed is slow but read speed is good.

Q: - Explain RAID 5?

RAID Level 5 is based on rotating parity with striping technique. RAID-5 stores parity information but not redundant data (but parity information can be used to reconstruct data). The storage capacity of Software RAID level 5 is equal to the capacity of the member partitions, minus the size of one of the partitions if they are of equal size. The performance of RAID 5 is based on parity calculation process but with modern CPUs that usually is not a very big problem. In RAID 5 read and write speeds are good.

Q: - Which kernel module is required for Software RAID?

“md” module

Q: - which utility or command is used for creating software RAID’s for RHEL5?

mdadm

Q: - Can we create software RAID during Linux installation?

Yes, we can create Software RAID during Linux Installation by “Disk Druid”

Q: - What is the role of chunk size for software RAID?

Chunk size is very important parameter on which RAID performance based.

We know stripes go across disk drives. But how big are the pieces of the stripe on each disk? The pieces a stripe is broken into are called chunks. To get good performance you must have a reasonable chunk size.

For big I/Os we required small chunks and for small I/Os we required big chunks.

Q: - What is SWAP Space?

Swap space in Linux is used when the amount of physical memory (RAM) is full. If the system needs more memory resources and the RAM is full, inactive pages in memory are moved to the swap space. While swap space can help machines with a small amount of RAM, it should not be considered a replacement for more RAM. Swap space is located on hard drives, which have a slower access time than physical memory.

Q: - What are the steps to create SWAP files or Partition?

- Create swap partition or file
- Write special signature using “mkswap”
- Activate swap space by “swapon -a” command
- Add swap entry into /etc/fstab file

Q: - How you will create swap file of size 4 GB and explain swap file entry in /etc/fstab file?

Use “dd” command to create swap file.
dd if=/dev/zero of=/SWAPFILE bs=1024 count=4
mkswap /SWAPFILE
swapon -a
Entry into /etc/fstab file.
/SWAPFILE swap swap defaults 0 0

Q: - Tell me the steps to remove the swap file?

Firstly disable the swap file by “swapoff” command.
Remove Swap file entry from /etc/fstab file.
Now remove the swap file by “rm” command.

Q: - What can we do with “parted” command or utility?

- View the existing partition table
- Add partitions from free space or additional hard drives
- Change the size of existing partitions

Q: - How we will check free space on drive /dev/sda with parted command?

```
#parted /dev/sda  
print
```

Q: - Can we resize the size of a partition?

Yes, we can resize the size of partition by “parted” command.

```
#parted /dev/sda  
print
```

To resize the partition, use the resize command followed by the minor number for the partition, the starting place in megabytes, and the end place in megabytes. For example:

```
resize 3 1024 2048
```

After resizing the partition, use the print command to confirm that the partition has been resized correctly, is the correct partition type, and is the correct file system type.

Q: - What is LVM?

LVM stands for Logical Volume Manager. LVM, is a storage management solution that allows administrators to divide hard drive space into physical volumes (PV), which can then be combined into logical volume groups (VG), which are then divided into logical volumes (LV) on which the filesystem and mount point are created.

Q: - What are the steps to create LVM?

- Create physical volumes by “pvcreate” command
#pvcreate /dev/sda2
- Add physical volume to volume group by “vgcreate” command
#vgcreate VLG0 /dev/sda2
- Create logical volume from volume group by “lvcreate” command.
#lvcreate -L 1G -n LVM1 VLG0

Now create file system on /dev/sda2 partition by “mke2fs” command.
#mke2fs -j /dev/VLG0/LVM1

Q: - What is the difference between LVM and RAID?

RAID provides redundancy but LVM doesn't provide Redundancy.

Q: - What are LVM1 and LVM2?

LVM1 and LVM2 are the versions of LVM.
LVM2 uses device mapper driver contained in 2.6 kernel version.
LVM 1 was included in the 2.4 series kernels.

Q: - What is Volume group (VG)?

The Volume Group is the highest level abstraction used within the LVM. It gathers together a collection of Logical Volumes and Physical Volumes into one administrative unit.

Q: - What is physical extent (PE)?

Each physical volume is divided chunks of data, known as physical extents; these extents have the same size as the logical extents for the volume group.

Q: - What is logical extent (LE)?

Each logical volume is split into chunks of data, known as logical extents. The extent size is the same for all logical volumes in the volume group.

Q: - Explain LVM snapshot?

LVM snapshots allow the administrator to create a new block device which presents an exact copy of a logical volume, frozen at some point in time.

Q: - How you will check on Your server or system device-mapper is installed or not?

Check the following file.
#cat /proc/misc
if this file contains "device-mapper" term it means device mapper is installed on your system.

Q: - How are snapshots in LVM2 different from LVM1?

In LVM2 snapshots are read/write by default, whereas in LVM1, snapshots were read only.

Q: - What is the maximum size of a single LV?

For 2.4 based kernels, the maximum LV size is 2TB.
For 32-bit CPUs on 2.6 kernels, the maximum LV size is 16TB.
For 64-bit CPUs on 2.6 kernels, the maximum LV size is 8EB.

Q: - If a volume group named as VG0 already exists but i need to extend this volume group up to 4GB.Explain all steps?

Firstly create Physical volume (/dev/sda7) of size 4GB.
Now run following command.
vgextend VG0 /dev/sda7

Q: - If a volume group VG0 have 3 PV's (/dev/sda6, /dev/sda7, /dev/sda8) but i want to remove /dev/sda7 pv from this VG0?

vgreduce VG0 /dev/sda7

Q: - Which command is used to extend a logical volume?

```
lvextend --size + /dev//  
resize2fs /dev//
```

Q: - Tell me all steps to remove a LVM?

To remove a logical volume from a volume group, first unmount it with the unmount command:

```
umount /dev//
```

and then use the lvremove command:

```
lvremove /dev//
```

Q: - Which command is used to create LVM Snapshot?

```
vcreate --size -s -n
```

The lvcreate command is used to create a new logical volume, meaning there must be free physical extents in the logical volume group to create a snapshot. The -s option means that the LV is a snapshot, is the name of the new LV created, and is the name of the LV from which to create the snapshot.

Q: - Is there any relation between modprobe.conf file and network devices?

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

For Example :-

```
[root@localhost ~]# cat /etc/modprobe.conf  
alias eth0 b44
```

Here b44 is the kernel module for network device eth0.

We can Confirm by following command (This module "b44" is present or not).

```
[root@localhost ~]# lsmod |grep b44  
b44 29005 0
```

Q: - What is the location of "network" file and what does this contains?

location :- /etc/sysconfig/network

This file contains following fields

```
NETWORKING=yes
```

```
NETWORKING_IPV6=no
```

```
HOSTNAME=localhost.localdomain
```

Q: - Which daemon is required to start Network services?

```
network
```

```
/etc/init.d/network start
```

Q: - What is the role of /etc/resolv.conf file?

In this file we sets the DNS servers (using their IP addresses) and the search domain. The values of the DNS servers are often added when the network is activated because the data can be provided by DHCP or a similar service.

Q: - What "neat" command will do?

neat command provides Graphical interface to change network settings for network devices.

Q: - Which protocol is required to allow local printing and print sharing?

Internet Printing Protocol (IPP) is required to allow local printing and print sharing.

Q: - What is CUPS?

CUPS stands for "Common UNIX Printing System". CUPS is a open source printing system developed by Apple Inc. CUPS uses the Internet Printing Protocol (IPP) to allow local printing and print sharing.

Q: - What is the location of log files for CUPS?

The log files for the CUPS printing system are located in the /var/log/cups/ directory.

Q: - What is YUM?

YUM stands for Yellow dog Updater, Modified because it is based on YUP, the Yellow dog Updater. Where does the name Yellow dog come from? Yellow Dog 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.

Q: - What are the advantages of YUM?

- Automatic resolution of software dependencies.
- Multiple software locations at one time.
- Ability to specify particular software versions or architectures.

Q: - How you will install software by YUM?

yum install

Q: - Which option is required to assume the answer "yes" to any questions asked during installation of package dependencies for YUM?

The "-y" option is used to assume the answer "yes".

For Example

yum -y install squid

Q: - How to remove a software by YUM?

yum remove

Q: - How Many Run Levels present in Linux?

There are 7 run levels, with each having its own properties.

- 0: Halt the system
- 1: Single-user mode
- 2: Not used
- 3: Multi-user mode with text login
- 4: Not used
- 5: Multi-user mode with graphical login
- 6: Reboot

Q: - Which configuration file is required to change the Run Level of Server or system?

/etc/inittab

To change the default run level, modify this line.

id:5:initdefault:

Q: - Explain architectures required for RPMs?

noarch Architecture-independent, can run on any architecture
i386 Generic build for a 32-bit x86 system
i586 Sometimes used when building kernels for older x86 processors
Intel® Pentium® II, Intel Pentium III, Intel Pentium 4, AMD Athlon, and
i686 AMD Duron systems (Most RPMs for these architectures are built using the i386 architecture, with the kernel for these architectures being built with the i686 for optimal performance.)
x86_64 64-bit processors such as AMD Athlon64, AMD Opteron, and Intel EM64T
ia64 Intel® Itanium
ppc 32-bit IBM® POWER, IBM eServer, pSeries®, and IBM eServer iSeries
s390x 64-bit IBM eServer System z

Q: - How to install Linux software's by RPM?

rpm -ivh test-1.0-1.i386.rpm

Q: - If a file associated with test-1.0-1.i386.rpm deleted, than How we will recover that file?

We can reinstall this rpm again.

Q: - If you are getting error "package is already installed" but you have to install package any how. what option you will use?

rpm -ivh test-1.0-1.i386.rpm

Preparing... ##### [100%] package test-1.0-1 is already installed

In this case you can use "--replacepks" option.

rpm -ivh --replacepks test-1.0-1.i386.rpm

Q: - Which options are required to upgrade a RPM?

Upgrading a package is similar to installing one. Type the following command at a shell prompt:

rpm -Uvh test-2.0-1.i386.rpm

Q: - In which directory RPM database stored?

/var/lib/rpm

Q: - Explain the command "rpm -qa"?

It will queries all currently installed packages.

Q: - Explain the command "rpm -qf "?

it queries the RPM database for which package owns . When specifying a file, specify the absolute path of the file.

Q: - How to verify all installed packages?

rpm -Va

Q: - How to verify the signature of an rpm?

rpm -K test-1.0-1.i386.rpm

Q: - How to list PCI Devices on your server or System?

use "lspci" command.

Q: - What is the role of "Kudzu"?

Kudzu is used to Detect new Hardware

Q: - What happens when you add a new device after installation?

The Kudzu program runs each time the system boots and performs a hardware probe. If new hardware is found, Kudzu attempts to map it to a kernel module. If successful, the information is saved, and the device is configured.

Q: - How to Enable ACLs for /home partition?

Add following entry in /etc/fstab

```
LABEL=/home /home ext3 acl 1 2
```

Now remount /home partition with acl option.

```
mount -t ext3 -o acl /dev/sda3 /home
```

Q: - How to View ACLs for a file(test_file)?

```
getfacl test_file
```

Q: - How to remove an ACL?

```
setfacl --remove-all </pkgname></pkgname></addsize>
```