

## Total Question 103

---

### QUESTION 1

Which two utilities can you use to set up a job to run at a specified time?

- A. at and crond
- B. atrun and crontab
- C. at and crontab
- D. atd and crond

Answer: C

Explanation: The 'at' command is used to execute commands at a specified time and optional date.

A cron job is a program or script scheduled at a specified time. The 'crontab' program is used to create user cron jobs.

Reference: <http://www.oreillynet.com/linux/cmd/a/at.html>

<http://www.oreillynet.com/linux/cmd/c/crontab.html>

Incorrect Answers:

A: The Cron daemon (crond) is the program that runs the cron job at the specified time. It is not used to set up a cron job.

B: Atrun is used to run jobs scheduled by the 'at' program. It is not used to set up a job to run at a specified time.

D: Atd is the 'at' daemon. Similar to the cron daemon, it is the program that runs the jobs scheduled with the 'at' command.

---

### QUESTION 2

After creating a backup of the users home directories called backup.cpio you are asked to restore a file called memo.ben. What command should you type?

Answer: `cpio -iF backup.cpio memo.ben`

Explanation: The 'cpio' command is used to create backups or restore files from a backup. The -i option is to extract something. The F option is to specify a file. Here we are extracting memo.ben from a file named backup.cpio.

Reference: <http://www.oreillynet.com/linux/cmd/c/cpio.html>

---

### QUESTION 3

You wish to restore the file memo.ben which was backed up in the tarfile MyBackup.tar. What command should you type?

Answer: `tar xf MyBackup.tar memo.ben`

Explanation: Tarfiles are created using the 'tar' utility. Therefore, you should use the 'tar' utility to extract the files. The x option is to extract and the f option is to specify a filename to extract from.

Reference: <http://www.oreillynet.com/linux/cmd/t/tar.html>

---

**QUESTION 4**

When is the most important time to restore a file from your backup?

- A. On a regular scheduled basis to verify that the data is available.
- B. When the system crashes.
- C. When a user inadvertently loses a file.
- D. When your boss asks to see how restoring a file works.

Answer: A

Explanation: According to 'best practice', you should regularly restore files to verify that your backup procedures are working. It's no good backing up files regularly if you are unable to restore files when needed.

Incorrect Answers:

B: When the system crashes, you may need to restore your whole system. However, this shouldn't be the only time you restore files.

C: When a user loses a file, it will need to be restored. However, you should prepare for this eventuality by regularly testing your backup/restore process to ensure you are able to restore a file when needed.

D: When your boss asks to see how restoring a file works, you want it to work when you show him. This is why you should test your backup/restore processes.

---

**QUESTION 5**

Which one of the following factors does not play a role in choosing the type of backup media to use?

- A. How frequently a file changes.
- B. How long you need to retain the backup.
- C. How much data needs to be backed up.
- D. How frequently the backed up data needs to be accessed.

Answer: A

Explanation: Your choice of backup media will depend on its capacity, its shelf life and the time it takes to access the data. The frequency of file changes is irrelevant.

Incorrect Answers:

B: Different backup media can be kept for varying periods of time. You should find out from the manufacturers how long a backup media can be kept without losing its data.

C: Obviously, your choice of backup media will depend on the amount of data to be backed up. For example, a CD-ROM can hold around 700MB of data while tapes can hold up to hundreds of gigabytes of data.

D: Your choice of backup media will also depend on the time it takes to retrieve data from the media. Reading data from a CD-ROM or DVD is much quicker than reading data from a tape.

---

**QUESTION 6**

You attempt to log out but receive an error message that you cannot. When you issue the jobs command, you see a process that is running in the background. How can you fix this so that you can

logout?

- A. Issue the kill command with the PID of each running command of the pipeline as an argument.
- B. Issue the kill command with the job number as an argument.
- C. Issue the kill command with the PID of the last command as an argument.
- D. Issue the kill command without any arguments.

Answer: C

Explanation: The kill command is used to send a signal to kill one or more process IDs. You must own the process or be a privileged user, otherwise the kill command will be ignored.

Reference: <http://www.oreillynet.com/linux/cmd/k/kill.html>

Incorrect Answers

- A: You need to end the background process. You know its process ID; therefore you should issue the process ID with the kill command to kill the appropriate process.
- B: You cannot use the job number with the kill command. You should use the process ID or process name.
- D: The kill command won't work if it doesn't know what process you want it to kill.

---

### QUESTION 7

The top utility can be used to change the priority of a running process? Another utility that can also be used to change priority is \_\_\_\_\_?

Answer: nice

Explanation: The 'nice' command enables you to run a command with a different priority. Nice -n <adjustment> command, increments the priority of 'command' by <adjustment>. You can increase the priority of a command by specifying a negative adjustment. For example, 'nice -n-5 command' will run 'command' with the priority increased by 5.

Reference: <http://www.oreillynet.com/linux/cmd/n/nice.html>

---

### QUESTION 8

You need to search the entire directory structure to locate a specific file. How could you do this and still be able to run other commands while the find command is still searching for you file?

- A. find / -name filename &
- B. find / -name filename
- C. bg find / -name filename
- D. &find / -name filename &

Answer: A

Explanation: The find command is used to locate files. / is the root directory, so searching from / will search the entire directory tree. The -name <filename> enables you to search for a file named <filename>. The ampersand character (&) is used to return control of the shell returning you to the command prompt, without have to wait for the command to execute.

Reference: <http://www.oreillynet.com/linux/cmd/f/find.html>

Incorrect Answers

B: With no ampersand (&) following the command, you will not be able to run other commands until the find command has completed its search.

C: The bg command is used to run a suspended job in the background if job control is enabled. However, the program or command would have to start and then be suspended for this to work.

D: The ampersand (&) must follow the command, not precede it.

---

**QUESTION 9**

In order to display the last five commands you have entered using the history command, you would type \_\_\_\_\_.

Answer: history 5

Explanation: The history command is used to display the previously entered commands. If you typed history with no arguments, you would likely get a long scrolling list of commands. By typing a number after 'history', you will display only the last <number> of commands.

Reference: <http://www.redhat.com/docs/manuals/linux/RHL-7.3-Manual/getting-started-guide/s1-q-and-ahistory-tips.html>

---

**QUESTION 10**

You telnet into several of your servers simultaneously. During the day, you sometimes get confused as to which telnet session is connected to which server.

Which of the following commands in your .profile would make it obvious to which server you are attached?

A. PS1='\h: \w>'

B. PS1='\s: \W>'

C. PS1='\!: \t>'

D. PS1='\a: \n>'

Answer: A

Explanation: The PS1 environment variable controls the prompt on the command line, and can be used by users to tell what system they are on, the directory they are currently in, the current date and more depending on how this variable is configured. The \h option is used to specify the hostname and the \w option will give the full path of the current working directory.

Reference: <http://ctdp.tripod.com/os/linux/tips/tipsps1.html>

Incorrect Answers:

B: The \s option is used to display the shell name. This won't give any indication of which machine you are connected to.

C: The \ option is used to display the history number of the current command. This won't give any indication of which machine you are connected to.

D: The \a option is used to display a new line. This won't give any indication of which machine you are connected to.

---

**QUESTION 11**

You have to type your name and title frequently throughout the day and would like to decrease the number of key strokes you use to type this. Which one of your configuration files would you edit to bind this information to one of the function keys?

Answer: `.inputrc`

Explanation: The `inputrc` file is used to map keystrokes to text or commands. You can use this file to make a function key display your name and title. Other common uses include mapping a function key to lock your computer or run a command.

Reference: <http://beyond.linuxfromscratch.org/view/cvs/postlfs/inputrc.html>

---

**QUESTION 12**

When typing at the command line, the default editor is the \_\_\_\_\_ library.

Answer: `readline`

Explanation: The default command line editor is the `Readline` library. As with most text editor programs, it allows certain keystrokes to aid in the writing/editing of a command. For example, there are keystroke combinations that allow you to jump to the beginning or end of the line, or to jump to the start or end of a previous word.

Reference: <http://www.cs.utah.edu/dept/old/texinfo/bash/rlman.html>

---

**QUESTION 13**

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

Answer: `echo $SHELL`

Explanation: The `'echo'` command is used to echo a string to standard output. `$shell` is an environment variable that reflects the current shell in use. Therefore, the `'echo $shell'` command will display the name and path of the shell you are using.

Reference: <http://www.santafe.edu/projects/echo/how-to/node30.html>

---

**QUESTION 14**

You have recently decided to convert from using a monolithic kernel to using a modular kernel. You have made the appropriate changes in your kernel configuration. Next you wish to compile your new kernel and modules and copy the modules to their proper location. What would you type to do this?

- A. `make modules modules_install`
- B. `make bzImage modules modules_install`
- C. `make mrproper modules modules_install`
- D. `make dep clean modules modules_install`
- E. `make dep clean bzImage modules modules_install`

Answer: E

Explanation: This command consists of multiple make commands on the same line:

The first part of the command, make dep, actually takes your configuration and builds the corresponding dependency tree. This process determines what gets compiled and what doesn't.

The next step, make clean, erase all previous traces of a compilation so as to avoid any mistakes in which version of a feature gets tied into the kernel.

The next step, make bzImage does the full compilation of the kernel.

The next two steps, make modules and make modules\_install will compile the modules and copy them to their appropriate location.

Reference: <http://www.openna.com/community/articles/security/v1.3-xml/chap7sec84.html>

Incorrect Answers

A: This command will compile the modules, but not the kernel.

B: You need the make dep command to build the dependency tree.

C: Make mrproper is similar to make clean except that it doesn't delete any binaries. However, there is no kernel image specified in this command.

D: There is no kernel image specified in this command.

---

### QUESTION 15

To allow a user to mount a CD and read from it, which entry should be put into /etc/fstab?

- A. /dev/cdrom /mnt/cdrom iso9660 noauto,user,ro 0 0
- B. /dev/cdrom /mnt/cdrom iso9660 noauto,uid=user,gid=group,ro 0 0
- C. /dev/cdrom /mnt/cdrom iso9660 noauto,User,ro 0 0
- D. /dev/cdrom /mnt/cdrom iso9660 noauto,usermap,ro 0 0
- E. /dev/cdrom /mnt/cdrom iso9660 noauto,owners,ro 0 0

Answer: A

Explanation: This entry in the fstab file allows any user to mount the CD-ROM (/dev/cdrom) in the /mnt/cdrom directory. Iso9660 is the file system for the CD-ROM. Noauto means that the CD-ROM won't be automatically mounted when the system boots. The first '0' means that the CD-ROM shouldn't be backed up and the second '0' means that the CD-ROM file system shouldn't be checked for errors when the machine boots.

Reference: Roderick W. Smith. Sybex Linux + Study Guide: Page 400/1.

Incorrect Answers:

B: The syntax of this entry is incorrect.

C: The 'user' field should be lowercase.

D: Usermap is an invalid entry for the user field.

E: Owners is an invalid entry for the user field.

---

### QUESTION 16

What is the usable disk space of a RAID 5 array of five 18 GB drives with one drive dedicated as a spare?

- A. 18GB
- B. 34GB
- C. 54GB

- D. 72GB
- E. 90GB

Answer: C

Explanation: The question states that you have 5 18GB drives, but one is dedicated as a spare. Therefore, you have the use of 4 drives which equals 72GB. When using RAID 5, parity data is written across the disks, using the equivalent of one disk's space (18GB). Therefore, the total useable space is  $72 - 18 = 54$ GB.

Reference: <http://www.pc.ibm.com/us/infobrf/raidfin.html>

Incorrect Answers

A: The total usable space is 54GB, not 18GB.

B: The total usable space is 54GB, not 34GB.

D: The equivalent of one drive is used for parity. Therefore, the total useable space is  $72 - 18 = 54$ GB, not 72GB.

E: One drive is spare and the equivalent of one drive is used for parity. Therefore, the total useable space is

---

#### **QUESTION 17**

You have to mount the /data filesystem from an NFS server (srv1) that does not support locking. Which of the following mount commands should you use?

- A. `mount -a -t nfs`
- B. `mount -o locking=off srv1:/data /mnt/data`
- C. `mount -o nolocking srv1:/data /mnt/data`
- D. `mount -o nolock srv1:/data /mnt/data`
- E. `mount -o nolock/data@srv1 /mnt/data`

Answer: D

Explanation: If you are mounting a volume that does not support locking, you need to use the nolock option with the mount command. The nolock option tells the system to not use the NFS locking protocol.

Reference: [http://docsrv.caldera.com:8457/cgi-bin/info2html?\(am-utils.info.gz\)opts%2520Option](http://docsrv.caldera.com:8457/cgi-bin/info2html?(am-utils.info.gz)opts%2520Option)

Incorrect Answers

A: This answer has the wrong command options.

B: 'Locking=off' is the wrong option. It should be 'nolock'.

C: 'Nolocking' is the wrong option. It should be 'nolock'.

E: /data@srv1 is the wrong syntax. It should be <servername>:/<foldername>.

---

#### **QUESTION 18**

To list the filesystem available from the NFS server 'castor', the command " \_\_\_\_\_ -e castor" can be used.

Answer: showmount

Explanation: The showmount is used to display information about NFS filesystems. The -e option is used

to specify an exported filesystem.

Reference: <http://www.oreillynet.com/linux/cmd/s/showmount.html>

---

**QUESTION 19**

You want to check what shares are offered by a Windows system. Which of the following commands could you use to perform this task?

- A. mmblookup
- B. showshares
- C. smbclient
- D. smbstatus
- E. listshares

Answer: C

Explanation: The smbclient command with the -L can be used to display the shares on a Windows system. The syntax is smbclient -L //<servername>.

Reference: <http://www.winnetmag.com/Articles/Index.cfm?ArticleID=8897>

Incorrect Answers

- A: Mmblookup is the incorrect command to display shares on a Windows system.
  - B: Showshares is the incorrect command to display shares on a Windows system.
  - D: Smbstatus is the incorrect command to display shares on a Windows system.
  - E: Listshares is the incorrect command to display shares on a Windows system.
- 

**QUESTION 20**

What file in the /proc file system will show you the parameters passed to your kernel at boot time?

- A. /proc/apm
- B. /proc/stat
- C. /proc/kmsg
- D. /proc/sys/kernel/sysrq
- E. /proc/cmdline

Answer: E

Explanation: The /proc/cmdline file contains parameters passed to the kernel at system boot time.

Reference: <http://www.dobit.com/emblin/embhowto.htm>

Incorrect Answers

- A: The kernel parameters are in the /proc/cmdline file, not the /proc/apm file.
  - B: The kernel parameters are in the /proc/cmdline file, not the /proc/stat file.
  - C: The kernel parameters are in the /proc/cmdline file, not the /proc/kmesg file.
  - D: The kernel parameters are in the /proc/cmdline file, not the /proc/sys/kernel/sysrq file.
- 

**QUESTION 21**

When an ext2 partition is formatted, a fixed percentage of the blocks on the disk are reserved for use by the root user. After the filesystem has been created this percentage can be modified using which



utility?

- A. tune2fs
- B. mke2fs
- C. e2fsck
- D. mount
- E. hdparm

Answer: A

Explanation: The tune2fs utility can be used to modify the reserved blocks. For example, the tune2fs - u <username> command can be used to allow a user to use the reserved blocks.

Reference: <http://www.oreillynet.com/linux/cmd/t/tune2fs.html>

Incorrect Answers:

- B: You need the tune2fs utility, not mke2fs.
- C: E2fsck is used to check a disk for bad blocks. It is not used for reserved blocks.
- D: The mount command is used to mount a filesystem. It is not used for reserved blocks.
- E: Hdparm is used for tuning a hard disk for performance. It is not used for reserved blocks.

---

### QUESTION 22

You are asked to provide access through your FTP server to a network share available from an NT server running on your local network- For this purpose, you will need \_\_\_\_\_ support in the kernel and to mount the NT share using the smbmount command line utility:

Answer: smbfs

Explanation: Windows NT uses SMB (Server Message Blocks) for network communications. In order to be able to use the smbmount command to mount a Windows NT share, your kernel must have smbfs (server message block file system) support.

Reference: <http://uranus.it.swin.edu.au/~jn/linux/smbfs/>

---

### QUESTION 23

On an ext2 filesystem, a running daemon has created a large logfile that is beginning to fill the disk. After deleting the file with an "rm-f" command as root, "df" shows that the space is still in use even though the file is not shown using "ls". To reclaim this space you must:

- A. Restart the daemon.
- B. Unmount and remount the filesystem.
- C. Use sync.
- D. Recreate the file.
- E. Run fsck.

Answer: A

Explanation: If you have a daemon which writes a log file and keeps that file open for writing then removing the file will not free up the disk space. The filesystem still sees the program as having a

reference to it. Therefore the filesystem will not free up that disk space. The only way to free the space is to restart the daemon

Reference: <http://mail.gnu.org/pipermail/bug-fileutils/2001-February/001495.html>

Incorrect Answers:

B: Unmounting and remounting the filesystem is unnecessary and may not free the space.

C: Sync is used to write the buffers to disk. It will not free the space.

D: Recreating the file will not free the space because the daemon has a reference to the old file.

E: Fsck is a file system checking tool. It won't free the space because it won't recognize it as corrupted.

---

### QUESTION 24

While checking the logfiles on your log server, you notice that all client machines are showing up by IP address rather than by hostname, although DNS is configured and running. How would you ensure that host entries show by name rather than by IP?

A. Restart named and then syslogd on the log server.

B. Add the central logging server to all inbound logging hosts' /etc/hosts.

C. Recompile syslogd to add remote logging support.

D. Restart syslogd on the inbound logging clients to force DNS resolution.

E. Add all inbound logging hosts to /etc/hosts on the log server, then restart syslogd.

Answer: E

Explanation: I don't know why the DNS resolution isn't working for the syslog daemon. It could be that there are lots of log entries and that the DNS requests are timing out. Therefore, adding the inbound logging hosts to /etc/hosts on the log server will enable local hostname resolution, thus negating the need to use DNS.

Incorrect Answers

A: The question states that DNS is configured and running and therefore does not need to be restarted.

B: This won't work. The clients are able to contact the logging server. Adding the central logging server to all inbound logging hosts' /etc/hosts files won't affect how the logging server records the log entries.

C: Remote logging support is already enabled because the IP addresses are being logged.

D: DNS resolution needs to be forced on the server, not the clients.

---

### QUESTION 25

You are trying to boot a system and change the root password, which you do not know. What do you type at the LILO prompt?

A. linux /etc/passwd

B. linux norootpass

C. linux disable passwords

D. linux init=/bin/bash

E. linux passwd=0

Answer: D

Explanation: If you forget the root password, you can boot init into the shell and change the password

using the following commands:

boot: Linux init=/bin/sh

bash# mount -o remount / -rw

bash# passwd root

Reference: Michael J. Tobler. New Riders, Inside Linux: Page 466.

Incorrect Answers

A: linux /etc/passwd is not a valid boot prompt command.

B: linux norootpass is not a valid boot prompt command.

C: linux disable passwords is not a valid boot prompt command.

E: linux passwd=0 is not a valid boot prompt command.

---

### QUESTION 26

You need to use grep to search for specific log entries. Given the following three log entries, which grep command will match only one line? Assume that every pattern matches at least one line.

Jun 16 01:46:18 hostname pumpd[10]: PUMP: got an offer

Jun 17 21:52:28 hostname kernel: SCSI subsystem driver Revision: 1.00

Jul 20 11:09:01 hostname /USR/SBIN/CRON[1800]: (mail) CMD runq

A. grep "hostname\ [^\]\*[A-Z]\*:"

B. grep "Ju[I-Z]\[0-9].\*:.1"

C. grep "hostname,\*[pumpd]\*[10]"

D. grep "[0-9]:[1-8]\*\ host.\*\(.\*\)"

E. grep "US\*[^]]]\*:"

Answer: D.

---

### QUESTION 27

How can you determine who has scheduled at jobs?

A. at -l

B. at -q

C. at -d

D. atwho

Answer: A

Explanation: The at -l command is the same as the atq command. It will list the user's pending jobs, unless the user is a privileged user; in which case, everybody's jobs are listed

Reference: <http://www.oreillynet.com/linux/cmd/a/at.html>

Incorrect Answers

B: The -q option is used to place the job in a specified queue. It does not display who has scheduled jobs.

C: The -d option is used to delete a specified job. It does not display who has scheduled jobs.

D: This is an invalid command.

---

### QUESTION 28

You want to create a compressed backup of the users home directories. What utility should you use?

Answer: tar

Explanation: The tar utility is used to archive multiple files into one 'tarball'. The -z option invokes another utility called gzip and instructs it to compress the files before tar archives them.

Reference: <http://www.oreillynnet.com/linux/cmd/t/tar.html>

---

**QUESTION 29**

You are covering for another system administrator and one of the users asks you to restore a file for him. You locate the correct tarfile by checking the backup log but you do not know how the directory structure was stored. What command can you use to determine this?

- A. tar fx tarfile dirname
- B. tar tvf tarfile filename
- C. tar ctf tarfile
- D. tar tvf tarfile

Answer: D

Explanation: You can list the contents of a 'tarball' with the tar tvf tarfile command. The t option is used to list the files and directories. The v option runs the command in verbose mode. The f option allows you to specify the name of the tarball (a tarball is a common name for an archive created with the tar utility) with the f <filename> option.

Reference: <http://www.oreillynnet.com/linux/cmd/t/tar.html>

Incorrect Answers

A: The syntax of this command is wrong. The x must come before the f. This also does not list the contents of the file.

B: This command would list the path to 'filename'. Although this would be required information to restore a file, the question states that you want to view the directory structure.

C: The c option is used to create a tarball which isn't required in this question.

---

**QUESTION 30**

The easiest, most basic form of backing up a file is to \_\_\_\_\_ it to another location.

Answer: copy

Explanation: The easiest way to backup a file is to copy it to another location. Having a backup copy of a file is always recommended.

---

**QUESTION 31**

When planning your backup strategy you need to consider how often you will perform a backup, how much time the backup takes and what media you will use.

What other factor must you consider when planning your backup strategy?

Answer: what to backup

Explanation: The first thing to consider when planning a backup strategy is what you are going to back up. Then you can think about the amount of data this will be. This will affect your other decisions such as what media to use etc.

---

**QUESTION 32**

What key combination can you press to suspend a running job and place it in the background?

Answer: ctrl-z

Explanation: You can suspend a currently running job by using the Ctrl + z keystroke. This will stop the job, but it won't end it. The job will be available to be resumed. Note: you can only stop jobs that were started in your current shell.

Reference: <http://unix.about.com/library/weekly/aa072301b.htm>

---

**QUESTION 33**

Using command substitution, how would you display the value of the present working directory?

- A. echo \$(pwd)
- B. echo pwd
- C. \$pwd
- D. pwd | echo

Answer: A

Explanation: The echo command can be used to display the contents of variables. The present working directory is held in the pwd variable. Echo \$(pwd) will display the contents of the pwd variable. Other commands that would work are echo \$ PWD and echo "\$PWD".

Reference: <http://www.bolthole.com/solaris/ksh-beforeyoustart.html>

Incorrect Answers:

- B: Echo pwd would display the text 'pwd'.
  - C: \$pwd doesn't do anything although \$PWD would work.
  - D: pwd | echo doesn't do anything.
- 

**QUESTION 34**

Every time you attempt to delete a file using the rm utility, the operating system prompts you for confirmation. You know that this is not the customary behavior for the rm command. What is wrong?

- A. rm has been aliased as rm -i
- B. The version of rm installed on your system is incorrect.
- C. This is the normal behavior of the newest version of rm.
- D. There is an incorrect link on your system.

Answer: A

Explanation: The -i option with the rm command runs the command in 'interactive' mode. This will cause rm to prompt you for the deletion of a file.

Reference: <http://www.oreillynet.com/linux/cmd/r/rm.html>

Incorrect Answers

B: The fact that rm is prompting for a confirmation indicates that the version of rm is compatible with your system.

C: This is not the normal behaviour for rm although it will prompt you if you are attempting to delete a write protected file.

D: The rm command is running the rm program so there is not an incorrect link.

---

### QUESTION 35

In your present working directory, you have the files.

maryletter

memo1

MyTelephoneandAddressBook

What is the fewest number of keys you can type to open the file MyTelephoneandAddressBook with vi?

- A. 6
- B. 28
- C. 25
- D. 4

Answer: A

Explanation: Tab completion is where you can type the first few letters of a command or filename then press tab to automatically complete the command or filename. You need to type enough letters so that there is only one command or filename starting with those letters. In this question you could type v then i then space then m then y then tab. This equals six keystrokes. There is only one filename starting with 'my' so this file will be opened.

Reference: [http://www.cmp.liv.ac.uk/misc/guide/linux\\_guide/node28.html](http://www.cmp.liv.ac.uk/misc/guide/linux_guide/node28.html)

Incorrect Answers:

B: You need a minimum of 6 keystrokes, not 28.

C: You need a minimum of 6 keystrokes, not 25.

D: You need a minimum of 6 keystrokes, not 4.

---

### QUESTION 36

After typing in a new command and pressing enter, you receive an error message indicating incorrect syntax. This error message originated from?

- A. The shell.
- B. The operating system.
- C. The command.
- D. The kernel.

Answer: C

Explanation: When you run a 'command' you are actually instructing the shell to run a program. If the

shell can find the program, it will run it. The shell knows how to start the program, but it doesn't know the syntax of the program/command. If you get an error saying 'incorrect syntax', the error will be coming from the program.

Incorrect Answers

A: The shell knows how to start the program, but it doesn't know the syntax of the program/command. A shell error message would be for example, '<command>: Command not found.'

B: The operating system runs the shell. It doesn't know about specific commands.

D: The kernel is effectively the operating system. It doesn't know about specific commands.

---

**QUESTION 37**

A single machine acts as a mail server, web server, and gateway to the Internet for the rest of your internal network. Why shouldn't you also use this machine as your central log host?

A. It may reduce web server performance.

B. The remote logging may have a negative impact on network performance.

C. If the web server crashed, log messages from other hosts would be lost.

D. Under high load, syslogd on the web server may start rejecting messages, and clients would try to log the error, creating a recursive loop between the clients and the log host.

E. If the security of your server is compromised, an attacker would have access to log information from all your hosts.

Answer: E

Explanation: You are running three services that connect directly to the Internet (mail server, web server and gateway). This in itself poses a security risk. The logs will contain a lot of information that an attacker would be able to access if the attacker gained access to your system.

Incorrect Answers

A: The logging may have a minimal effect on the web server. This is not as much a problem as the security risk in answer E.

B: It is unlikely that the remote logging will have any negative impact on the network performance.

C: If the web server crashed you may not be able to access the logs. However, this is not as much a problem as the security risk in answer E.

D: This just wouldn't happen.

---

**QUESTION 38**

Which of the following parameters in your smb.conf file specifies the relationship between Windows/SMB usernames and Linux/UNIX usernames?

A. smb usernames = /etc/smbusers

B. username map = /etc/smbusers

C. map usernames = lowercase nospace

D. smb usernames = map to same unix name

E. usernames map = /etc/windows-usernames.map

Answer: B

Explanation: To map Windows usernames to Linux/UNIX usernames, you would specify the username map = /etc/smbusers parameter in the smb.conf file. You would then add mappings for each user account in the smb.conf file. For example, windows\_username = linux\_username.

Reference: <http://www.mandrakeuser.org/docs/connect/csamba2.html>

Incorrect Answers

A: smb usernames = /etc/smbusers is the wrong syntax for this parameter.

C: map usernames = lowercase nospace is the wrong syntax for this parameter.

D: smb usernames = map to same unix name is the wrong syntax for this parameter.

E: usernames map = /etc/windows-usernames.map is the wrong syntax for this parameter.

---

**QUESTION 39**

Assuming modules for all supported filesystems have been loaded, which file contains a list of filesystems that can be currently mounted on the system?

A. /proc/filesystems

B. /dev/filesystems

C. /etc/filesystems

D. /var/fs/filesystems

E. /etc/config/filesystems

Answer: A

Explanation: In the file /proc/filesystems you can find which filesystems your kernel currently supports. (If you need a currently unsupported one, you'll need to insert the corresponding module or recompile the kernel.)

Reference: <http://www.fifi.org/cgi-bin/man2html/usr/share/man/man5/fs.5.gz>

Incorrect Answers

B: The list of currently supported filesystems is in the filesystems file in the /proc directory, not the /dev directory.

C: The list of currently supported filesystems is in the filesystems file in the /proc directory, not the /etc directory.

D: The list of currently supported filesystems is in the filesystems file in the /proc directory, not the /var/fs directory.

E: The list of currently supported filesystems is in the filesystems file in the /proc directory, not the /etc/config directory.

---

**QUESTION 40**

An ext2 filesystem is used by an application that frequently reads a large number of small files. Performance can be improved by mounting the filesystem with the \_\_\_\_\_ option.

A. atime

B. noatime

C. noexec

D. nosuid

E. sync



Answer: B

Explanation: Linux records information about when files were created and last modified as well as when it was last accessed. There is a cost associated with recording the last access time. The ext2 file system of Linux has an attribute that allows the super-user to mark individual files such that their last access time is not recorded. This may lead to significant performance improvements on often accessed frequently changing files.

Reference: <http://www.tldp.org/LDP/solrhe/Securing-Optimizing-Linux-RH-Edition-v1.3/chap6sec73.html>

Incorrect Answers

A: The atime option will record the last access time for each file which may degrade system performance.

C: Noexec is a mount flag to not allow any executables to be run from the filesystem. This won't work since the files are likely to be text files.

D: Nosuid is a mount flag to disallow any setuid binaries on the filesystem. This will not improve system performance.

E: The sync command is used to write the buffers to disk. This will not improve system performance.

---

### QUESTION 41

You decide to use the logical volume manager (LVM) to manage four 4GB disk drives. After creating the volume group, how would you create a 10GB logical volume called big-app?

- A. `vgcreate -p 10g -n /dev/vg01/big-app`
- B. `vgcreate -l 2560 /dev/vg01/big-app`
- C. `mkldm -v 10240 -n /dev/vg01/big-app`
- D. `lvcreate -v 10240 /dev/vg01/big-app`
- E. `lvcreate -l 2560 vg01 -n big-app`

Answer: E

Explanation: When you create a volume group, it will have a physical extent size of 4MB by default, unless otherwise specified. When you add disks to the volume group, the disk space is divided into chunks equal to the physical extent size (4MB by default). When you create a logical volume with the `lvcreate` command, the `-l` option is used to specify the size of the logical drive in 'logical extents'. The logical extents are the same size as the physical extents. Therefore, to create a 10GB logical drive, you would specify 2560 logical extents (2560 x 4MB = 10GB). `Vg01` is the name of the volume group in which to create the logical volume. The `-n` option allows you to enter a name for the logical volume. In this case 'big-app'.

Reference: [http://devresource.hp.com/STKLI/man/11iv1.5/lvcreate\\_1m.html](http://devresource.hp.com/STKLI/man/11iv1.5/lvcreate_1m.html)

Incorrect Answers

A: The `vgcreate` command is used to create the volume group.

B: The `vgcreate` command is used to create the volume group.

C: `Mkldm` is an invalid command.

D: `-v` is an incorrect option to create a logical volume.

---

### QUESTION 42

You maintain daily backups of a large file, as well as calculating an MD5 checksum with `md5sum`. When verifying the contents of one such backup, you notice that the new checksum is different from the previous one by only one byte. What does this tell you about the contents of the file?

- A. A single character in the original file has been modified.
- B. 1/32nd of the original file has been modified.
- C. 1/128th of the original file has been modified.
- D. It tells you that the original file has been modified.
- E. The contents of the file are in reverse order from the original.

Answer: D

Explanation: The MD5 (Message Digest number 5) value for a file is a 128-bit value similar to a checksum. This value is calculated according to the contents of a file. If a file has changed, the MD5/checksum value will be different.

Reference: <http://www.iay.pwp.blueyonder.co.uk/threel/tech/tools/md5.htm>

Incorrect Answers

- A: A different checksum means that the file has changed. However, it offers no indication of how much of the file has changed.
- B: A different checksum means that the file has changed. However, it offers no indication of how much of the file has changed.
- C: A different checksum means that the file has changed. However, it offers no indication of how much of the file has changed.
- E: A different checksum means that the file has changed. However, it offers no indication of how much of the file has changed.

---

### **QUESTION 43**

You are creating a script with demands that the previous command execute correctly. How would you correctly test the exit status of the previous command in BASH?

- A. if [ "\$#" -eq "0" ]; then...
- B. if [ "\$?" -eq "0" ]; then...
- C. if [ '\$#' == 0 ]; then...
- D. if [ '\$?' == '0' ]; then...
- E. if [ \$@ -eq 0 ]; then...

Answer: B

Explanation: The variable "\$?" checks the exit status of the last command run. The -eq "0" statement is used to check whether a condition is true. The statement if [ "\$?" -eq "0" ]; then... will check that the last command executed correctly and run the next part of the script.

Reference: <http://www.bolthole.com/solaris/ksh-basics.html>

Incorrect Answers

- A: The variable is "\$?" not "\$#".
  - C: The variable is "\$?" not "\$#".
  - D: The variable is "\$?" not '\$?' (double quotes, not single quotes).
  - E: The variable is "\$?" not \$@.
-

**QUESTION 44**

You are having problems with programs crashing on an SMP system, and would like to run your system in non-SMP mode for troubleshooting purposes. What is the correct parameter to pass to the kernel at boot time to force it to use a single CPU?

- A. block-cpu-1
- B. cpucount=1
- C. disable-cpu
- D. nosmp
- E. enable\_smp=no

Answer: D.

Explanation: The nosmp option can be used at boot time to disable SMP (symmetric multiprocessing), thus causing the system to run in uniprocessor (single processor) mode.

Incorrect Answers

A: This is an incorrect option.

B: Cpucount=1 will run one processor, but the single processor will run in SMP mode. Cpucount=0 will run no processors in SMP mode effectively disabling SMP.

C: Disable-cpu is an incorrect option.

E: enable\_smp=no is an incorrect option.

---

**QUESTION 45**

You want to export a local file system /data, and permit read-write access for all users on host

A. In

addition, the root account on hostA should be allowed root level access. All other hosts are to have

- A. /data hostA(rw,no\_root\_squash) (ro)
- B. /data hostA(allow\_root) -ro
- C. /data -ro,rw=hostA(root)
- D. /data hostA=rw,root \*=ro
- E. /data hostA(rw,all\_squash,anonid=0) @all(ro)

Answer: A

Explanation: The format of the /etc/exports lines is 'directoryname hostname(options)'. In this case you are exporting the /data directory. HostA has the (rw,no\_root\_squash) permissions applied and everyone else gets the (ro) permission. The rw permission allows HostA read/write permissions. The no\_root\_squash option gives the root account on HostA root access to the /data directory. Everyone else gets the ro permission which means read only.

Reference: [http://www.comptechdoc.org/os/linux/usersguide/linux\\_ug nfs.html](http://www.comptechdoc.org/os/linux/usersguide/linux_ug nfs.html)

Incorrect Answers

B: The option to allow root access is no\_root\_squash, not allow\_root.

C: The syntax in this answer is incorrect.

D: The syntax in this answer is incorrect.

E: The syntax in this answer is incorrect.

---

**QUESTION 46**

The \_\_\_\_\_ command can be used to change the default root device hardcoded into a kernel image.

Answer: rdev

Explanation: The rdev command is used to change the default root device hardcoded into a kernel image.

Reference: <http://www.linuxcentral.com/linux/man-pages/rdev.8.html>

---

**QUESTION 47**

Several of your users have been scheduling large at jobs to run during peak load times. How can you prevent anyone from scheduling an at job?

- A. Delete the file /etc/at.deny
- B. Create an empty file called /etc/at.deny
- C. Create two empty files: /etc/at.deny and /etc/at.allow
- D. Create an empty file called /etc/at.allow

Answer: D.

Explanation: The /etc/at.allow and the /etc/at.deny files are used to control who is allowed to run the 'at' command. If the file /etc/at.allow exists, only usernames mentioned in it are allowed to use the 'at' command, and the /etc/at.deny file is ignored.

Reference: <http://ccrma-www.stanford.edu/planetccrma/man/man5/at.deny.5.html>

Incorrect Answers

A: The /etc/at.allow file is read before the /etc/at.deny file. If an /etc/at.allow file exists, any names in that file will be able to use the 'at' command. Deleting the /etc/at.deny file may work, but only if no /etc/at.allow file exists.

B: An empty file called /etc/at.deny is the default on a Linux system and allows anyone to use the 'at' command.

C: Creating two empty files: /etc/at.deny and /etc/at.allow would also work because an empty /etc/at.allow file would prevent the use of the 'at' command. However, it is unnecessary to create both files.

---

**QUESTION 48**

You want to create a compressed backup of the users home directories so you can issue the command `gzip /home/* backup.gz` but it fails. The reason it failed is that gzip will only compress one \_\_\_\_ at a time.

Answer: file

Explanation: The command: `gzip <filename> backup.gz` will compress <filename> and rename it to backup.gz. This only works with a single file. To compress multiple files into one file (archive), you should use the tar utility with the z option. Tar can archive multiple files into a single file (archive). The z option causes tar to use gzip to compress the files first.

Reference: <http://www.oreillynet.com/linux/cmd/g/gzip.html>

---

**QUESTION 49**

You need to view the contents of the tarfile called MyBackup.tar. What command would you use?

Answer: tar tf MyBackup.tar

Explanation: You can list the contents of a 'tarball' with the 'tar tf tarfile' command. The t option is used to list the files and directories. The f option allows you to specify the name of the tarball (a tarball is a common name for an archive created with the tar utility) with the f <filename> option.

Reference: <http://www.oreillynet.com/linux/cmd/t/tar.html>

---

**QUESTION 50**

When you back up only the files that have changed since the last backup, this is called a \_\_\_\_\_ backup.

- A. Partial
- B. Differential
- C. Full
- D. Copy

Answer: B

Explanation: When you run a full backup, the files are marked as having been backed up (the archive attribute is cleared). When a file is created or changed, it is marked as 'not backed up' (the archive attribute is set). A differential backup backs up any files created or changed since the last full backup (the files marked as 'not backed up'). It does not mark files as having been backed up (in other words, the archive attribute is not cleared).

Reference: [http://www.raid-unix-mac-disk-daterecovery-service.com/diferential\\_backup.htm](http://www.raid-unix-mac-disk-daterecovery-service.com/diferential_backup.htm)

Incorrect Answers

- A: A partial backup is not an actual backup type. It is just a description of backing up a few selected files.
  - C: A full backup backs up all files and marks them as having been backed up by clearing the archive attribute.
  - D: A copy backup backs up all files but does not mark them as having been backed up.
- 

**QUESTION 51**

You have been given the job of administering a new server. It houses a database used by the sales people. This information is changed frequently and is not duplicated anywhere else. What should you do to ensure that this information is not lost?

- A. Create a backup strategy that includes backing up this information at least daily.
- B. Prepare a proposal to purchase a backup server.
- C. Recommend that the server be made part of a cluster.
- D. Install an additional hard drive in the server.

Answer: A

Explanation: To ensure that data isn't lost, it should be backed up. The question states that the information

is changed frequently, so it should be backed up frequently.

Incorrect Answers

B: A backup server usually runs backup software. This may not be necessary and is no use without a backup strategy.

C: Clustering the server would require additional servers and would be very expensive. Furthermore, depending on the type of cluster, you may only have one set of hard disks containing the information.

D: Installing an additional hard drive would only work if the data was regularly backed up to the additional hard drive. However, if the server failed, the data would still be unavailable.

---

**QUESTION 52**

What utility can you use to show a dynamic listing of running processes?

Answer: top

Explanation: The 'top' command is used to provide information (frequently refreshed) about the most information isn't dynamically refreshed.

Reference: <http://www.oreillynet.com/linux/cmd/t/top.html>

---

**QUESTION 53**

You previously ran the find command to locate a particular file. You want to run that command again. What would be the quickest way to do this?

A. `fc -l find <enter>`

n

B. `history -l find <enter>`

history n

C. Retype the command

D. `fc -n find`

Answer: A

Explanation: The -l option used with the fc command is used to list the commands saved in the 'history'. The 'fc -l find' command will display all recent commands starting with the word 'find'. After pressing enter, the list is displayed and you can recall the command by entering the number (n) of the command.

Reference: <http://www.computerhope.com/unix/uhistory.htm>

Incorrect Answers

B: The syntax of the 'history' command is wrong.

C: Whether it would be quicker to retype the command or not depends on the name of file you were looking for previously. It is unlikely that this is a trick question, so the answer would be to use the fc command.

D: The 'fc -n find' command would display the recent 'find' commands, but without the command numbers. It would not run the required command.

---

**QUESTION 54**

Which of the following environment variables determines your working directory at the completion of a successful login?

- A. HOME
- B. BASH\_ENV
- C. PWD
- D. BLENDERDIR

Answer: A

Explanation: The HOME environment variable determines your working directory when you log on. This is typically /home/<username> for a normal user account or the root directory (/) for the root user. The HOME environment variable also determines the directory you will be taken to if you enter the 'cd' command with no arguments.

Reference: <http://www.isu.edu/departments/comcom/unix/workshop/environment.html>

Incorrect Answers

B: The BASH\_ENV variable is used for non-interactive shells. It does not determine your working directory when you log on.

C: The PWD variable contains the current working directory. It does not determine your working directory when you log on. The 'pwd' command is used to display the full path to your current directory.

D: The BLENDERDIR variable is used with a piece of software named 'Blender'. It does not determine your working directory when you log on.

---

### QUESTION 55

After experimenting with vi as your command line editor, you decide that you want to have vi your default editor every time you log in. What would be the appropriate way to do this?

- A. Change the /etc/inputrc file
- B. Change the /etc/profile file
- C. Change the ~/.inputrc file
- D. Change the ~/.profile file

Answer: C

Explanation: The .inputrc file is used to control your shell. You can set keystrokes to perform specified functions with this file. Another setting that can be changed is the command line editor. If you want this setting to only affect you, you edit the file in your home directory (~/. signifies your home directory).

Reference: [http://ctdp.tripod.com/os/linux/howlinuxworks/linux\\_hlkeyprogs.html](http://ctdp.tripod.com/os/linux/howlinuxworks/linux_hlkeyprogs.html)

Incorrect Answers

A: The settings in the /etc/inputrc file would affect all users. If you want this setting to only affect you, you edit the file in your home directory.

B: The default editor is not set in this file.

D: The default editor is not set in this file.

---

### QUESTION 56

You want to enter a series of commands from the command line. What would be the quickest way to do this?

- A. Press enter after entering each command and its arguments.

- B. Put them in a script and execute the script.
- C. Separate each command with a semi-colon (;) and press enter after the last command.
- D. Separate each command with a / and press enter after the last command.

Answer: C

Explanation: You can enter multiple commands on one line by separating them with a semi-colon (;). Pressing enter after the last command will run the commands.

Reference: <http://unix.about.com/library/tips/bltip016.htm>

Incorrect Answers

- A: Pressing enter after a command will run the command before you can enter another command.
- B: A script listing the commands would work but this isn't the quickest way of doing it.
- D: You need to use semi-colons (;) to separate the commands, not forward slashes (/).

---

### QUESTION 57

You have compiled and installed a new kernel on your SCSI based machine. After installing the new kernel, the boot process stops at a point with the error "VFS PANIC: Unable to mount root FS." You can boot again off the old kernel without any problems. Given that /etc/modules.conf is correct and that the SCSI controller is selected as a module in the kernel, what most likely is the cause?

- A. The module failed to build.
- B. The new kernel can't initialize the SCSI controller.
- C. There is no initrd image for the new kernel.
- D. SCSI disk support isn't enabled in the kernel.
- E. SCSI generic support isn't enabled in the kernel.

Answer: C

Explanation: The question states that the machine is SCSI based and you can boot to the old kernel. This indicates that the system is successfully booting from the SCSI drive (when using the old kernel). The SCSI controller module needs to be loaded at boot time before the system is able to mount the root file system. To load the SCSI controller module at boot time, you need an initrd image for the new kernel.

Reference: <http://www.linuxhelp.co.za/RedHat61/rhref/s1-sysadmin-build-kernel.htm#S2-SYSADMININITRD>

Incorrect Answers

- A: It is unlikely that the module failed to build.
- B: The new kernel can't initialize the SCSI controller. However, the reason for this is most likely to be that there is no initrd image for the new kernel.
- D: SCSI support can be loaded as a module if an initrd image exists. It does not have to be enabled (compiled) in the kernel.
- E: SCSI support can be loaded as a module if an initrd image exists. It does not have to be enabled (compiled) in the kernel.

---

### QUESTION 58

The 'user' option in /etc/fstab allows a normal user to mount/unmount filesystems. When used on removable devices, this can allow unaudited applications to be made available on your system. For security reasons, you may wish to disable:



- The suid bit.
- Device nodes.
- Running of executables.
- Writing to the mounted filesystem.

Which of the following is a valid /etc/fstab entry which implements AT LEAST one of these features?

- A. /dev/fd0 /mnt ext2 ro,user,noauto,noexec,nodev 0 0
- B. /dev/cdrom /mnt iso9660 rw,user,nobin,nosuid,nodev 00
- C. /dev/cdrom /mnt iso9660 ro|user|!dev|!suid|!bin 0 0
- D. /dev/fd0 /mnt vfat rw+user+noexec+nodev+nosuid 0 0
- E. /dev/cdrom /mnt auto ro|user|!auto

Answer: A

Explanation: The 'ro' option means read only. This means that the drive can only be mounted in read only mode and therefore, cannot be written to. The noexec option prevents the running of executable files.

Reference: [http://www.humbug.org.au/talks/fstab/fstab\\_options.html](http://www.humbug.org.au/talks/fstab/fstab_options.html)

Incorrect Answers

B: The rw option will allow the drive to be mounted in read/write mode. The question states that writing to the filesystem should be disabled.

C: The options must be separated by commas. Therefore the syntax in this answer is incorrect.

D: The options must be separated by commas. Therefore the syntax in this answer is incorrect.

E: The options must be separated by commas. Therefore the syntax in this answer is incorrect.

---

### QUESTION 59

You have been asked by your management to come up with a backup solution that covers not only data loss, but also situations where the entire system, or building, is destroyed. Your solution should also protect against data theft. Which of the following plans provides the most secure redundant backup and storage solution?

A. Once a week, all of your systems receive a full system backup to tape.

Those tapes are stored in a secured location in your facility.

B. Once a week, all of your systems receive an incremental system backup to tape.

Those tapes are stored in a secured location in remote facility.

C. Every night, all of your systems receive an incremental system backup to tape.

Those tapes are stored in a secured location in a remote facility.

D. Every night, all of your systems receive an incremental system backup to tape, and once a month, all systems receive a full backup to tape.

Those tapes are stored in a secured remote facility.

E. Once a week, all of your systems receive a full backup to tape.

Twice a month, all of your systems receive a full backup to CD.

The tapes are stored in a secured remote facility.

The CD's are stored locally.

Answer: D

Explanation: Your backup strategy should allow for backups to occur as often as possible. Daily backups are recommended but this isn't always possible because the amount of time a full backup would take. For this reason, daily incremental backups are recommended. An incremental backup will only backup the files that have been created or changed since the last backup, thus saving time and backup media space. You should however perform full backups on a regular basis (monthly is recommended). This will allow the easy restoration of a failed system.

The backup tapes should be stored in a secured remote facility, in case the building is destroyed.

Incorrect Answers

A: If you back up the files once a week, you could lose up to one week's data. For this reason, daily backups are recommended.

B: If you back up the files once a week, you could lose up to one week's data. For this reason, daily backups are recommended.

C: You should perform full backups on a regular basis (monthly is recommended). This will allow the easy restoration of a failed system.

E: If you back up the files once a week, you could lose up to one week's data. For this reason, daily backups are recommended.

---

**QUESTION 60**

Which of the following daemons must be running on an NFS server?

- A. portmap
- B. nfsiod
- C. nfsd
- D. xinetd
- E. mountd

Answer: A, C, E.

Explanation: If you want to provide NFS service to other hosts, you have to run the `rpc.nfsd` and `rpc.mountd` daemons on your machine. As RPC-based programs, they are not managed by `inetd`, but are started up at boot time and register themselves with the portmapper; therefore, you have to make sure to start them only after `rpc.portmap` is running.

Portmap provides port information to clients requesting RPC services on the server. Mountd determines which filesystem and devices are available to which machine and users. Nfsd is the daemon on the server that handles client filesystem requests.

Reference: <http://www.linuxvalley.it/encyclopedia/ldp/guide/nag2/x-087-2-nfs.daemons.html>

Incorrect Answers

B: Nfsiod runs on an NFS client machine to service asynchronous I/O requests to its server. It improves performance but is not required for correct operation.

D: Xinetd is a replacement for `inetd`, the internet services daemon, and offers improved functionality. However, it is not a requirement to run NFS.

---

**QUESTION 61**

Which process had the Process ID number 1?

- A. bash

- B. kernel
- C. init
- D. it varies
- E. none

Answer: C

Explanation: As with files, all processes that run on a GNU/Linux system are organized in the form of a tree. The root of this tree is init. Each process has a number (its PID, Process ID), together with the number of its parent process (PPID, Parent Process ID). The PID of init is 1, and so is its PPID: init is its own father.

Reference: <http://www.mandrakeuser.org/docs/mdoc/ref/process.html>

Incorrect Answers

- A: Process ID number 1 represents init, not bash.
- B: Process ID number 1 represents init, not the kernel.
- D: Init always has a process ID of 1. It does not vary.
- E: Init has a process ID of 1.

---

### QUESTION 62

For a change to the primary Samba configuration file smb.conf to take effect, it is necessary to:

- A. Restart the smbd and nmbd processes.
- B. Send a HUP signal to the smbd and nmbd processes.
- C. Do nothing.
- D. Reboot the system.
- E. Restart the Samba subsystem.

Answer: A

Explanation: Whenever you make changes to the smb.conf file, it is necessary to restart the smbd and nmbd processes. Smbd regularly reads the smb.conf file and implements any changes. However, these changes don't affect any previously established connections. To apply the changes to any previously established connections, you must restart smbd and nmbd.

Reference: [http://us2.samba.org/samba/ftp/cvs\\_current/packaging/SGI/relnotes.html](http://us2.samba.org/samba/ftp/cvs_current/packaging/SGI/relnotes.html)

Incorrect Answers

- B: If you have Samba configured to be started by inetd, you could send inetd a HUP signal to restart it, but you wouldn't send smbd and nmbd a HUP signal.
- C: To apply the changes to any existing connections, it is necessary to restart smbd and nmbd.
- D: Rebooting the system would work if you have configured samba to start automatically. However, restarting the entire system is unnecessary.
- E: You should restart nmbd as well as smbd.

---

### QUESTION 63

After installing a package using dpkg, you find that the package manager tools no longer function. You isolate the problem to a broken library and you have a copy of the fixed library in a Debian .deb file. How can you extract files from a .deb file without using the Debian package manager?

- A. deb packages are compressed tar files with custom scripts.  
Use GNU 'tar' to extract the file.
- B. deb packages are redhat (rpm) packages with different fields.  
Use 'rpm' to extract the file.
- C. deb packages are simply gzipped cpio files.  
Use 'gunzip' to decompress the package and then use 'cpio' to extract the file.
- D. deb packages use a proprietary format and the file cannot be extracted without specialized tools.
- E. deb packages are ar archives with a special magic number.  
Use 'ar' to extract the data member and then use GNU 'tar' to extract the file.

Answer: E

Explanation: Debian archive (.deb) files can be parsed and manipulated by the utility ar. The precise contents of Debian archive files changed since Debian 0.93. The new contents are understood by versions of the primary package tool, dpkg, later than 0.93.76, and is described in the "deb"(5) man page. The old format is described in "deb-old"(5).

Using the command `ar -t foo_VVV-RRR.deb`, you'll see that a Debian archive file contains these members:

- `debian-binary`: Contains one or more lines; currently it contains only one line giving the version number (2.0) of the Debian package format.
- `control.tar.gz`: A compressed (gzip'd) tar file which contains the Debian control files for this package. (Confusingly, one of these files, and the only one which is required, is itself named `control`.)
- `data.tar.gz`: A compressed (gzip'd) tar file which contains the executables, libraries, documentation, etc., associated with this package. In other words, this component is the filesystem data part of a Debian package.

You can extract files from the .tar.gz files using the 'tar' utility.

Reference: <http://flits102-126.flits.rug.nl/~erik/debian/debian-faq-6.html>

Incorrect Answers

- A: You must first use the 'ar' utility to open the .deb file. Then you can use 'tar' to extract the required files.
- B: .deb files are not rpm (redhat package manager) files, and therefore cannot be opened with the rpm utility.
- C: .deb packages are not gzipped cpio files, and therefore cannot be opened with gunzip and cpio.
- D: .deb packages can be opened with the 'ar' utility; therefore, specialist tools are not required.

---

## QUESTION 64

You are in charge of a domain. Your developers have asked that mirrors of certain sites be placed as actual directories off the default path. Specifically they have asked that the `ftp.example-debian.org` Debian tree should be mapped at `/usr/local`. Assume that `ftp.example-debian.org` does an NFS export of their site. What would be the correct entry in the `/etc/auto.master` file?

- A. `/usr/local/debian ro ftp.example-debian.org:/pub/debian`
- B. `/usr/local/debian /etc/auto.debian` with `/etc/auto.debian` containing `debian-ro,soft,intr:ftp.exampledebian.org:/pub/debian`

- C. /usr/local/debian :etc/auto.debian with /etc/auto.debian containing debian:rw,soft,intr:ftp.exampledebian::/pub/debian
- D. /etc/auto.debian with /etc/auto.debian containing debian-ro,soft,intr:ftp.exampledebian.org:/pub/debian
- E. /etc/auto.debian with /etc/auto.debian containing debian:rw,soft,intr:ftp.exampledebian.org:/pub/debian

Answer: B

Explanation: Autofs uses the automount daemon to manage your mount points by only mounting them dynamically when they are accessed. Autofs consults the master map configuration file /etc/auto.master to determine which mount points are defined. It then starts an automount process with the appropriate parameters for each mount point. Each line in the master map defines a mount point and a separate map file that defines the filesystems to be mounted under this mount point.

In this question, the /etc/auto.master file would contain the line "/usr/local/debian /etc/auto.debian". /usr/local/debian is the mount point on the local machine. /etc/auto.debian is the name of the map file that defines what should be mounted at the mount point.

The /etc/auto.debian file should contain "debian-ro,soft,intr:ftp.example-debian.org:/pub/debian". This contains the mount point (debian), followed by some mount options (ro,soft,intr) followed by the directory to be mounted in the form of hostname:directory (ftp.example-debian.org:/pub/debian).

Reference: <http://www.redhat.com/docs/manuals/linux/RHL-7.3-Manual/custom-guide/s1-nfs-mount.html>

Incorrect Answers

- A: There is no map file defined in this line.
- C: The syntax is wrong. usr/local/debian :etc/auto.debian should be /usr/local/debian /etc/auto.debian.
- D: There is no mount point (usr/local/debian) defined in this answer.
- E: There is no mount point (usr/local/debian) defined in this answer.

---

### QUESTION 65

A dumb terminal on the serial line /dev/ttyE0 is losing characters when receiving large blocks of data from the server. Suspecting a flow control problem, you wish to examine the complete list of settings for this line on the server. Please type the one command that completes this command line:

\_\_\_\_\_ -a </dev/ttyE0

Answer: stty

Explanation: The stty command works on the current terminal by default, but by using the input redirection ("<") feature of the shell, we can have stty manipulate any tty device. The -a option is used to display all configuration settings.

Reference: <http://www.oreillynet.com/linux/cmd/s/stty.html>  
<http://www.tldp.org/LDP/nag2/x-087-2-serial-configuration.html>

---

### QUESTION 66

You have written a script called users to parse the passwd file and create a list of usernames. You want to have this run at 5 am tomorrow so you can see the results when you get to work. Which of the following commands will work?

- A. at 5:00 wed usrs
- B. at 5:00 wed -b usrs
- C. at 5:00 wed -l usrs
- D. at 5:00 wed -d usrs

Answer: A

Explanation: The 'at' command is used to execute commands at a specified time and optional date. It can contain an optional date, formed as a month and date, a day of the week, or a special keyword (today or tomorrow). An increment can also be specified.

Reference: <http://www.oreillynet.com/linux/cmd/a/at.html>

Incorrect Answers

B: No options are required to run the script at the specified time. Furthermore, the options should be specified before the time and date, not after them.

C: No options are required to run the script at the specified time. Furthermore, the options should be specified before the time and date, not after them.

D: No options are required to run the script at the specified time. Furthermore, the options should be specified before the time and date, not after them.

---

#### **QUESTION 67**

You need to copy all the files and directories contained in the home directory to another location. What utility can you use for this?

- A. cpio
- B. cp
- C. mv
- D. mvdir

Answer: B.

Explanation: The 'cp' command is used to copy files or directories from one location to another. The -r option makes the command recursive which means it will copy an entire directory structure from one location to another.

Reference: <http://squat.net/pusci/doc/pxii-cursus/copy-mv.html>

Incorrect Answers

A: The cpio command can be used to copy all the files into a single archive file at another location. However, it would be easier to copy the contents of the /home directory with the cp command. The cpio command is often used to create tape backups of Linux systems.

C: The 'mv' command is used to move files, not copy them.

D: The 'mvdir' command is used to move directories, not copy them.

---

#### **QUESTION 68**

When you only back up one partition, this is called a \_\_\_\_\_ backup.

- A. Differential
- B. Full

- C. Partial
- D. Copy

Answer: C

Explanation: If you are backing up just one partition and not the entire system, you are backing up only 'part' of the system. This is therefore a partial backup.

Incorrect Answers

A: A differential backup backs up all files that have been changed or created since the last full or incremental backup. A differential backup does not describe what portion of a system was backed up.

B: A full backup backs up all the specified files. A full backup does not describe what portion of a system was backed up.

D: A copy backup backs up all the specified files. The difference between a full backup and a copy backup is that a full backup marks the files as having been backed up, whereas a copy backup doesn't. A copy backup does not describe what portion of a system was backed up.

---

### QUESTION 69

You issue the command jobs and receive the following output:

[1]- Stopped (tty output) pine

[2]+ Stopped (tty output) MyScript

How would you bring the MyScript process to the foreground?

- A. fg %2
- B. ctrl-c
- C. fg MyScript
- D. ctrl-z

Answer: A

Explanation: You can bring a background job to the foreground by executing the "fg" command. If there are several background jobs, then you must indicate which job you wish to move to the foreground by indicating its job number. The syntax would be fg %<job number>.

Reference: [http://www.itworld.com/nl/lnx\\_tip/10052001/](http://www.itworld.com/nl/lnx_tip/10052001/)

Incorrect Answers

B: Ctrl-c is used to stop a running command/job. It does not bring a background job to the foreground.

C: If there are several background jobs, then you must indicate which job you wish to move to the foreground by indicating its job number.

D: Ctrl-z is used to suspend a running command/job without ending it. It does not bring a background job to the foreground.

---

### QUESTION 70

In order to display the last five commands you have entered using the fc command, you would type \_\_\_\_\_.

Answer: fc -l -5

Explanation: The -l option used with the fc commands is used to list the previously entered commands. You can specify how many commands to list with the -<number> option after the -l option, for example, fc -l -5.

Reference: [http://sseti.udg.es/marga/books/O'Reilly-The\\_Linux\\_Web\\_Server-CDBookshelfv1.0/linux\\_web/lnut/ch07\\_06.htm](http://sseti.udg.es/marga/books/O'Reilly-The_Linux_Web_Server-CDBookshelfv1.0/linux_web/lnut/ch07_06.htm)

---

### QUESTION 71

A variable that you can name and assign a value to is called a \_\_\_\_\_ variable.

Answer: user

Explanation: A user variable can be created and named and have a value assigned to it from the command line or from a script.

Reference: <http://ddart.net/shell/bourneshell/sh2a.html#2.1>

---

### QUESTION 72

You are entering a long, complex command line and you reach the right side of your screen before you have finished typing. You want to finish typing the necessary commands but have the display wrap around to the left. Which of the following key combinations would achieve this?

- A. Esc, /, Enter
- B. \, Enter
- C. ctrl-d, enter
- D. esc, /, ctrl-d

Answer: B

Explanation: There is a way to enter a long command such that it will be broken at the end of the top line and continued on the next. This can be accomplished by typing a backslash (\) character before pressing enter at the breakpoint, as follows:

```
$ echo This is a long command so why not break it here \  
> and start on the next line. <enter>
```

which gives as output:

```
This is a long command so why not break it here and start on the next line.
```

The > is the shell's way of letting the user know that the current line is a continuation of the previous line.

Reference: <http://pneuma.phys.ualberta.ca/~gingrich/research/shells/node13.html>

Incorrect Answers

- A: This key combination will not wrap the text.
  - C: This key combination will not wrap the text.
  - D: This key combination will not wrap the text.
- 

### QUESTION 73

You have elected to use the automounter and the autofs script. Your /etc/auto.master file contains the following:

```
/home /etc/auto.home  
/project /etc/auto.project
```



/data yp:data.map

If you change the contents of /etc/auto.project to include a new source path what must be done to access the new path?

- A. Shutdown and restart the local NFS client daemons.
- B. Run fsck on the affected mount point.
- C. Issue the /etc/init.d/autofs reload command.
- D. Add the newly mapped path to /etc/fstab.
- E. Nothing, simply access the newly mapped resource.

Answer: E

Explanation: Autofs uses the automount daemon to manage your mount points by only mounting them dynamically when they are accessed. Autofs consults the master map configuration file /etc/auto.master to determine which mount points are defined. It then starts an automount process with the appropriate parameters for each mount point. Each line in the master map defines a mount point and a separate map file that defines the filesystems to be mounted under this mount point. For example, the /etc/auto.misc file might define mount points in the /misc directory; this relationship would be defined in the /etc/auto.master file. Each entry in auto.master has three fields. The first field is the mount point. The second field is the location of the map file, and the third field is optional. The third field can contain information such as a timeout value. If the source path changes, you can access the new path without changing anything because the mount points are mounted and unmounted dynamically when accessed or closed.

Reference: <http://www.redhat.com/docs/manuals/linux/RHL-7.3-Manual/custom-guide/s1-nfs-mount.html>

Incorrect Answers

- A: As the mount points are mounted and unmounted dynamically, there is no need to restart the NFS daemons.
- B: It is not necessary to run fsck (file system checker) on the mount point.
- C: As the mount points are mounted and unmounted dynamically, there is no need to restart the autofs daemon.
- D: Fstab and autofs are two separate methods for mounting filesystems.

---

### **QUESTION 74**

While attempting to boot your i386 system, the boot process fails with a message that the root filesystem could not be mounted. Which stage of the boot process is failing?

- A. Kernel
- B. Boot loader
- C. BIOS POST
- D. Fsck
- E. Init

Answer: A

Explanation: When the kernel is loaded, it mounts the root filesystem. The root filesystem cannot be mounted; therefore the boot process has failed at the kernel stage.

Reference: <http://barclay.its.monash.edu.au/~kim/boot/boot.html>

Incorrect Answers

B: The boot loader is stored in the master boot record by default. The boot loader loads the kernel, which then mounts the root filesystem.

C: The BIOS POST (power on self check) is the first part of the boot process. It does not mount the root filesystem.

D: Fsck runs after the root filesystem is mounted.

E: Init runs after the root filesystem is mounted.

---

**QUESTION 75**

While installing a new Ethernet card you notice ifconfig is showing an odd IRQ for the device. What command will make lspci show which IRQ the card is actually using as seen by the PCI bus instead of as seen by the kernel?

A. `lspci -v -v`

B. `lspci -v -M`

C. `lspci -v -b`

D. `lspci -vM`

E. `lspci -m -v`

Answer: C.

Explanation: Lspci is a utility for displaying information about all PCI buses in the system and all devices connected to them. The -v option tells lspci to be verbose and display detailed information about all devices. The -b option is for bus-centric view. Show all IRQ numbers and addresses as seen by the cards on the PCI bus instead of as seen by the kernel.

Reference: <http://ccrma-www.stanford.edu/planetccrma/man/man8/lspci.8.html>

Incorrect Answers

A: You can't specify the same option twice (-v -v).

B: The -M option invokes bus mapping mode which scans the bus extensively to find all devices including those behind misconfigured bridges etc.

D: The options should be separated by hyphens (-). Furthermore, M is the wrong option.

E: The -m option is used to dump PCI device data in machine readable form for easy parsing by scripts.

---

**QUESTION 76**

You have added a new file system to /etc/exports, but users complain that they still get "Permission denied" errors when they try to mount the new file system. Which of the following is the best solution to this problem?

A. Reboot the server.

B. Add the option (no\_root\_squash) to the entry already added.

C. Restart NFS.

D. Run the command `exportfs -a`

E. Run the command `export -nfs`.

Answer: D

Explanation: The `exportfs` command makes local directories available for Network File System (NFS) clients to mount. This command is normally invoked during system startup by the `/etc/rc.nfs` file and uses information in the `/etc/exports` file to export one or more directories, which must be specified with full path names.

The `/etc/xtab` file lists directories that are currently exported. To display this file, enter the `exportfs` command without flags or arguments. To alter the file or to alter the characteristics of one of its directories, root users can edit the `/etc/exports` file and run the `exportfs` command. The `-a` option exports all directories listed in the `/etc/exports` file. Such alterations can be done at any time. Never edit the `/etc/xtab` file directly.

Reference: <http://www.unet.univie.ac.at/aix/cmds/aixcmds2/exportfs.htm>

Incorrect Answers

A: It is not necessary to reboot the server. The `exportfs` command can be run at any time.

B: The `no_root_squash` option is used to allow root permission to an exported directory to the root user on a remote machine.

C: It is not necessary to restart NFS.

E: `Export -nfs` is the wrong command.

---

### QUESTION 77

You have downloaded the patch from 2.2.18 to 2.2.19, and applied it to `/usr/src/linux`, where you had previously configured and built kernel 2.2.18. How would you avoid going through the entire configuration process again, and only configure options which are new to the patched kernel?

- A. `make oldconfig`
- B. `make reconfigure`
- C. `sh scripts/reconfig`
- D. Edit `.config` by hand

Answer: A

Explanation: The '`make oldconfig`' command will take the kernel config file named `".config"` and only ask the configuration questions which are not already answered in that file. This will avoid having to go through the entire kernel configuration process again.

Reference: [http://www.linuxchix.org/content/courses/kernel\\_hacking/lesson2](http://www.linuxchix.org/content/courses/kernel_hacking/lesson2)

Incorrect Answers

B: `Make reconfigure` is not a valid command.

C: `Sh scripts/reconfig` is not a valid command.

D: Editing the `.config` file by hand would be a difficult and very risky way of configuring the kernel.

---

### QUESTION 78

On a system with separate partitions for `/`, `/usr`, `/var`, `/tmp`, which filesystem[s] can safely be mounted read-only?

- A. `/var`, `/usr`
- B. `/var`
- C. `/usr`, `/`, `/tmp`
- D. `/usr`
- E. `/tmp`

Answer: D

Explanation: The /usr partition contains common executables and documents such as man pages (help files), all of which should not be changed by users. Therefore, this partition should be mounted as read only.  
Reference: <http://www.redhat.com/docs/manuals/linux/RHL-6.0-Manual/install-guide/manual/doc084.html>

Incorrect Answers

A: The /var (variable) partition is used for data that frequently changes such as log files and therefore cannot be read only.

B: The /var (variable) partition is used for data that frequently changes such as log files and therefore cannot be read only.

C: The /tmp (temporary) partition is used to store temporary files and therefore cannot be read only.

E: The /tmp (temporary) partition is used to store temporary files and therefore cannot be read only.

---

**QUESTION 79**

You routinely compress old log files. You now need to examine a log from two months ago. In order to view its contents without first having to decompress it, use the \_\_\_\_\_ utility.

Answer: zlibc

Explanation: Zlibc is a program that allows existing applications to read compressed (GNU gzip'ed) files as if they were not compressed. Zlibc will transparently uncompresses the data from these files as soon as they are read, just as a compressed filesystem would do.

Reference: <http://zlibc.linux.lu/zlibc.html>

---

**QUESTION 80**

As a system administrator, you are instructed to backup all the users home directories. Which of the following commands would accomplish this?

A. tar rf usersbkup home/\*

B. tar cf usersbkup home/\*

C. tar cbf usersbkup home/\*

D. tar rvf usersbkup home/\*

Answer: B

Explanation: The c option used with the tar command is used to create an archive. The f <filename> option allows you to specify a filename.

Reference: <http://www.oreillynet.com/linux/cmd/t/tar.html>

Incorrect Answers

A: The r option is used to append the files to an existing archive.

C: The b option is used to specify a block size. As no block size is specified, this answer is incorrect.

D: The r option is used to append the files to an existing archive.

---

**QUESTION 81**

What would be displayed as the result of issuing the command ps ef?

- A. A listing of the users running processes formatted as a tree.
- B. A listing of the stopped processes.
- C. A listing of all the running processes formatted as a tree.
- D. A listing of all system processes formatted as a tree.

Answer: A

Explanation: Without any options, the ps command displays the running processes associated with the current user ID. The e option displays the processes' environment and the f option displays the processes in a tree format, illustrating the relationship between parent and child relationships.

Reference: <http://www.oreillynet.com/linux/cmd/p/ps.html>

Incorrect Answers

- B: Only the running processes are listed, not the stopped processes.
- C: You would need the a option to display all the running processes.
- D: Only the user processes are listed, not the system processes.

---

### QUESTION 82

You have installed a new application but when you type in the command to start it you get the error message:

Command not found

What do you need to do to fix this problem?

- A. Add the directory containing the application to your path.
- B. Specify the directory's name whenever you run the application.
- C. Verify that the execute permission has been applied to the command.
- D. Give everyone read, write and execute permissions to the application's directory.

Answer: A

Explanation: One important environment variable is PATH, a list of directories separated by colons (':'). These directories are searched through to find commands. If you try to invoke command 'foo', all the directories in PATH (in that order) are searched for an executable file 'foo' (one with x-bit on). If a file is found, it is executed.

Reference: <http://www.tldp.org/HOWTO/mini/Path-3.html>

Incorrect Answers

- B: It is not necessary to specify the directory's name, if the directory is in the path.
- C: If you didn't have execute permission, you would get a permission denied error.
- D: It is not necessary to give everyone these permissions. Users shouldn't have write access to an application directory.

---

### QUESTION 83

You typed the following at the command line: `ls -al /home/ hadden`

What key strokes would you enter to remove the space between the '/' and 'hadden' without having to retype the entire line?

- A. Ctrl-B, Del
- B. Esc-b, Del
- C. Esc-Del, Del
- D. Ctrl-b, Del

Answer: B

Explanation: The Esc-b keystroke combination will move the cursor back one word (to the start of the word 'hadden'). The Del keystroke will delete the previous character; in this case, it will delete the space before the word 'hadden'.

Reference: [http://sseti.udg.es/marga/books/O'Reilly-The\\_Linux\\_Web\\_Server-CDBookshelfv1.0/linux\\_web/lnut/ch07\\_06.htm](http://sseti.udg.es/marga/books/O'Reilly-The_Linux_Web_Server-CDBookshelfv1.0/linux_web/lnut/ch07_06.htm)

Incorrect Answers

- A: The Ctrl-B keystroke will move the cursor back one letter.
- C: The Esc-Del keystroke will cut the previous word, for pasting later.
- D: The Ctrl-b keystroke will move the cursor back one letter. (Ctrl-b is the same as Ctrl-B).

---

#### **QUESTION 84**

What file will show you the IRQs being used by different hardware devices?

- A. /proc/interrupts
- B. /proc/irqs
- C. /proc/irq
- D. /proc/int
- E. /proc/ints

Answer: A

Explanation: The IRQs being used by the hardware devices are listed in the /proc/interrupts file.

Reference: [http://linuxcommand.org/man\\_pages/lsdev8.html](http://linuxcommand.org/man_pages/lsdev8.html)

Incorrect Answers

- B: The IRQs are not listed in the /proc/irqs file.
- C: The IRQs are not listed in the /proc/irq file.
- D: The IRQs are not listed in the /proc/int file.
- E: The IRQs are not listed in the /proc/ints file.

---

#### **QUESTION 85**

When setting up a client to log to a central logging server, you should:

- A. Start the syslogd daemon on the server with all of the clients in its hostlist.
- B. Add @servername to the appropriate log line in /etc/syslog.conf
- C. Use the server as a DHCP server for the client.
- D. Share the log file on the server using NFS.
- E. None of the above.

Answer: B

Explanation: The file `/etc/syslog.conf` contains information used by the system log daemon, `syslogd` to forward a system message to appropriate log files and/or users. When forwarding messages to a remote logging server, you would specify the name of a remote host, prefixed with an `@`, as with: `@server`, which indicates that messages are to be forwarded to the `syslogd` on the named host.

Reference: <http://www.unidata.ucar.edu/cgi-bin/man-cgi?syslog.conf+4>

Incorrect Answers

B: The `syslogd` daemon on the local machine needs to be configured to send messages to the remote logging server.

C: DHCP is unrelated to logging and is therefore not required.

D: The log file does not need to be shared using NFS.

E: You must edit the `syslog.conf` file; therefore, this answer is incorrect.

---

### QUESTION 86

You have been asked to block network access to an NFS sever. You need to block all access except NFS access. Which of the following actions would you take to achieve this?

A. Make sure that `xinetd` is switched off.

B. Place "ALL: ALL" in `/etc/hosts.deny` and "NFS: ALL" in `/etc/hosts.allow`

C. Add IPChains rules to deny all incoming packets except for portmapper

D. Place "ALL: ALL" in `/etc/hosts.deny` and "portmap: ALL" in `/etc/hosts.allow`

E. Ensure that the `nfs-access.o` module is configured into the kernel and use the command "`nfs-ctlallow <your IP range>`" to provide the required access

Answer: D

Explanation: The `hosts.allow` file is read before the `hosts.deny` file. This means that you can block access to 'all' in the `hosts.deny` file, but allow access to specific ports by specific hosts in the `hosts.allow` file. In this answer, we are blocking all ports to all hosts in the `hosts.deny` file. However, we are allowing access to the portmap service for all hosts in the `hosts.allow` file. (The portmap service is for access to NFS).

Reference: <http://www.mandrakeuser.org/docs/connect/cnfs2.html>

Incorrect Answers

A: `Xinetd` must be running.

B: NFS uses the portmapper service. Therefore, you should enter 'portmap: ALL' in the `hosts.allow` file.

C: IPChains is a firewall program. This may work (if you have IPChains running), however using the `hosts.allow` and `hosts.deny` files is much simpler.

E: The module and command in this question don't exist or are incorrectly named.

---

### QUESTION 87

What command would you type to use the `cpio` command to create a backup called `backup.cpio` of all the users home directories?

Answer: `find /home | cpio -o > backup.cpio`

Explanation: The `cpio` command expects to receive a list that contains one file per line. That is exactly the type of list that the `find` utility creates. The `ls` utility can also create this type of list, meaning that you will

see either of the ls or the find utilities used in conjunction with cpio. And since cpio archives a list of files it receives from standard input, you usually use a pipe (|) whenever you create an archive with the cpio utility. A lot of documentation suggests using the -print option with the find command. For example, find /home -print | cpio -o > backup.cpio. However, this is not required on Linux systems, and other systems that use GNU find, although it is required on Unix systems.

Reference: <http://www.unet.univie.ac.at/aix/cmds/aixcmds1/cpio.htm>

---

**QUESTION 88**

What is wrong with the following command?

```
tar cvfb //dev/tape 20
```

- A. You cannot use the c option with the b option.
- B. The correct command should be tar -cvfb /dev/tape20.
- C. The arguments are not in the same order as the corresponding modifiers.
- D. The files to be backed up have not been specified.

Answer: C

Explanation: The command should read tar cvfb /dev/tape 20 / . The letters c, v, f and b are the 'modifiers'. The arguments are the options for the modifiers and should be in the same order as the modifiers. The c modifier is to create an archive. The v modifier is for verbose mode. The f modifier specifies the name of the tar file and so needs an 'argument' (in this case a tape drive called /dev/tape). The b modifier is used to set a block size and so needs an 'argument' (in this case 20). Note that the arguments following the modifiers are in the same order as the modifier. The "f" precedes the "b" modifier so the arguments have the device before the blocking factor. The arguments must be in the same order as the modifiers, which can sometimes cause a little confusion.

After the modifiers and arguments have been entered, you need to enter the files to be backed up (in this case the root directory '/').

Reference: <http://freebooks.boom.ru/view/LinuxUnleashed/ch45/759-762.html>

Incorrect Answers

A: You can use the c option with the b option.

B: You don't need a hyphen (-) when specifying tar options. Furthermore, the files to be backed up haven't been specified.

D: The files to be backed up have been entered (the root partition '/'), but they are entered in the wrong place.

---

**QUESTION 89**

Many factors are taken into account when planning a backup strategy. The one most important one is how often does the file \_\_\_\_\_.

Answer: change

Explanation: The frequency of a file changing will determine the frequency of your backup. If the file changes often, you will need to back up the file often, otherwise the backed up version of the file will be an old version.

---



**QUESTION 90**

You enter the command  
cat MyFile | sort > DirList &  
and the operating system displays  
[4] 3499  
What does this mean?

- A. This is job number 4 and the PID of the sort command is 3499.
- B. This is job number 4 and the PID of the job is 3499.
- C. This is job number 3499 and the PID of the cat command is 4.
- D. This is job number 4 and the PID of the cat command is 3499.

Answer: A

---

**QUESTION 91**

In order to create a file called DirContents containing the contents of the /etc directory you would type \_\_\_\_\_.

Answer: ls /etc >DirContents

Explanation: Mostly all commands send their output to the screen or take input from the keyboard, but in Linux it is possible to send output to a file or to read input from a file. For example, the ls command sends it's output to screen; to send the output to a file, you can use the command ls > filename. This will send the output of the ls command to filename.

In this question, the ls command lists the contents of the /etc directory and sends the list to a file named DirContents.

Reference: <http://www.netti.hu/doc/LinuxShellScript/rpf.htm>

---

**QUESTION 92**

You are running out of space in your home directory. While looking for files to delete or compress you find a large file called .bash\_history and delete it. A few days later, it is back and as large as before. What do you need to do to ensure that its size is smaller?

- A. Set the HISTFILESIZE variable to a smaller number.
- B. Set the HISTSIZE to a smaller number.
- C. Set the NOHISTFILE variable to true.
- D. Set the HISTAPPEND variable to true.

Answer: A

Explanation: The .bash\_history file is a file in a user's home directory that contains a list of the (recent) commands issued by this user at the bash command line. This file can grow to up to the number of lines specified in the HISTFILESIZE variable; therefore, to reduce the maximum size of the file, you should set the HISTFILESIZE variable to a smaller number.

Reference: <http://www.slug-vt.org/bash.html>

Incorrect Answers

B: The HISTSIZE variable contains the number of commands in the history. When a user logs off, the commands are written to the bash\_history file.

C: There is no NOHISTFILE variable.

D: If the HISTAPPEND variable is set to true, the history will be appended to the history file (bash\_history), otherwise the file will be overwritten. Therefore, setting the HISTAPPEND variable to false, not true would work.

---

**QUESTION 93**

You would like to temporarily change your command line editor to be vi. What command should you type to change it?

Answer: set -o vi

Explanation: The readline support in the bash shell defaults to emacs editing mode. You can easily switch that to vi mode by issuing the following command: set -o vi. This will last until you logoff. The next time you log on, the default editing mode will be used.

Reference: <http://www.portico.org/index.php3?catList=26>

---

**QUESTION 94**

Which of the following interprets your actions when typing at the command line for the operating system?

- A. Utility
- B. Application
- C. Shell
- D. Command

Answer: C

Explanation: The "shell" is another name for the command shell or command interpreter. This is the program that gives you a command prompt, accepts the commands you type there, and basically makes the computer do what you tell it to.

The shell's job is to interpret your commands and run the programs you request. Linux was designed to be a multitasking operating system, which means you can run more than one program at one time. Linux was also designed as a multi-user OS, which means that you can have more than one shell running at the same time. (Each user gets his own shell at login.) As a user, you have access only to the programs you are running, not the ones other users are running (though you can run your own copy of the same program). The programs are kept separate because they are "enclosed" in a "shell".

Reference: <http://www.control-escape.com/lx-shell.html>

Incorrect Answers

A: A utility is a program that can be run from the shell.

B: An application is another name for a utility or program.

D: A command is what you enter to run a utility/program/application.

---

**QUESTION 95**

You have created a local ext2 file system on the third partition of your first IDE disk drive. You want

to facilitate easy manual mounting but you DO NOT wish the filesystem to be automatically mounted at a boot. What is the correct /etc/fstab entry?

- A. /dev/hda3/newfilesystem ext2 noboot 0 1
- B. /newfilesystem /dev/hda3 ext2 defaults 0 1
- C. /newfilesystem ext2 /dev/hda3 user 0 1
- D. /dev/hda3/newfilesystem ext2 noauto 0 1
- E. /dev/hda3 ext2 /newfilesystem defaults 0 -1

Answer: D

Explanation: /dev/hda3 indicates the 3rd partition on the first IDE hard disk (hda). Ext2 indicates the filesystem type. Noauto means that the filesystem will not be automatically mounted. The first '0' means that the filesystem shouldn't be backed up and the 1 means that the filesystem should be checked for errors when the machine boots.

Reference: [http://www.humbug.org.au/talks/fstab/fstab\\_options.html](http://www.humbug.org.au/talks/fstab/fstab_options.html)

Incorrect Answers

A: Noboot is an incorrect option.

B: The syntax of this command (the path of the filesystem) is incorrect.

C: The syntax of this command (the path of the filesystem) is incorrect.

E: The defaults option will use the default fstab options. The default is to automatically mount the filesystem at boot time.

---

### QUESTION 96

The command \_\_\_\_\_ flushes the filesystem buffers and ensures that the changes that you have made to a file are written to disk.

Answer: sync

Explanation: The sync command is used to flush the filesystem buffers and ensures that the changes that you have made to a file are written to disk. When you shutdown the system, the system runs the sync command to flush the buffers to disk, but the command can be manually entered at any time.

Reference: <http://www.oreillynet.com/linux/cmd/s/sync.html>

---

### QUESTION 97

How would you find out the version of the kernel in /usr/src/linux?

- A. cat /usr/src/linux/.version
- B. cat /usr/src/linux/VERSION
- C. Look in the README
- D. head -4 /usr/src/linux/Makefile

Answer: D

Explanation: The head command is used to display the first few lines of a file. The default is 10 lines but you can specify a number (in this case 4). The makefile is a script that tells the make utility how to build a

program or programs (in this case, the kernel). Most makefiles contain comments at the top of the file which describe the program and version information.

Reference: <http://www.opussoftware.com/tutorial/TutMakefile.htm>

Incorrect Answers

A: .version doesn't usually exist as a subdirectory of file. Rather, it is usually a directory /usr/src/linux.version or /usr/src/linux[VERSION] containing the kernel source.

B: VERSION doesn't usually exist as a subdirectory of file. Rather, it is usually a directory /usr/src/linux.version or /usr/src/linux[VERSION] containing the kernel source.

C: There usually isn't a README file containing version information.

---

### QUESTION 98

You suspect malicious behavior by one of your console session users. Which of the following methods could be used so that you will be notified whenever the suspect user is logged in? The method should not tip off the suspect user or affect overall system integrity or performance to a noticeable degree.

- A. Pipe the btmp file to a filter and launch a notification script if the user logs on.
- B. Insert into the suspect user's profile a script to notify you.
- C. Configure syslogd to pipe all auth log messages to a script which checks for the suspect user and then notifies you via email.
- D. Modify the user's login script to inform you of his presence and then exec itself with the real shell.

Answer: C

Explanation: Syslogd (the system log daemon) can be configured via the syslog.conf file. This file specifies where log entries should be written. You can configure syslogd to send authentication log messages to a script which checks for the suspect user and then notifies you via email.

Incorrect Answers

A: The btmp file is used to record failed logon attempts. This won't work because the user is able to log on successfully.

B: Inserting a script into the users profile file won't work because the user may notice the script if he/she looks at the profile file.

D: Modifying the users login script won't work because the user may notice the modification if he/she looks at the script.

---

### QUESTION 99

You want to compile a kernel with an experimental change that is distributed in "patch" format. However, you want to make sure that the patch works correctly before changing the original kernel source code. How can you test the patch before actually applying it?

- A. patch -p1
- B. patch --context
- C. patch --unified
- D. patch --dry-run

Answer: D

Explanation: The patch -dry-run command is used to test a patch before applying it. This will produce a text output listing all the files that would be patched. If there are no 'Failed' messages, then the patch is safe to install.

Reference: <http://www.hmug.org/man/1/patch.html>

Incorrect Answers

A: Patch -p1 will apply the patch.

B: Patch --context will apply the patch as a 'context diff' file.

C: Patch --unified will apply the patch as a 'unified diff' file.

---

### QUESTION 100

On a running system, where can you find specific information about the partition tables, such as major and minor device numbers, and number of blocks?

A. /proc/partitions

B. /proc/cpuinfo

C. /proc/fstab

D. /etc/partitions

E. /etc/fstab

Answer: A

Explanation: The /proc/partitions file contains information about the partition tables, such as major and minor device numbers, and number of blocks.

Reference: [http://www.goavatar.com/linux\\_pcmcia.htm](http://www.goavatar.com/linux_pcmcia.htm)

Incorrect Answers

B: The /proc/cpuinfo file contains information about the CPU, not the disks.

C: The /proc/fstab file contains information about mounted filesystems and permissions on the filesystems. It does not contain information such as the number of blocks.

D: The partitions file is in the /proc directory, not /etc.

E: The fstab file is in the /proc directory, not /etc.

---

### QUESTION 101

On bootup, LILO prints out LIL and stops. What is the cause of this?

A. The descriptor table is bad.

B. LILO failed to load the second stage loader.

C. LILO failed to load the primary stage loader.

D. LILO failed to locate the kernel image.

Answer: A

Explanation: If you only see LIL- when booting the system, it means that LILO could not load the map file (descriptor table).

Reference: Roderick W. Smith. Sybex Linux + Study Guide: Page 634.

Incorrect Answers

B: If LILO failed to load the second stage loader, you would see only LI during bootup.

- C: If LILO failed to load the primary stage loader, you wouldn't see anything during bootup.  
D: If LILO failed to locate the kernel image, you would see LILO during bootup followed by an error message about the missing kernel image.
- 

**QUESTION 102**

What file will tell you at what frequency the system processor is running?

- A. /proc/cpuinfo
- B. /proc/frequency
- C. /proc/speed
- D. /proc/mhz
- E. /proc/bogomips

Answer: A

Explanation: The /proc/cpuinfo file contains information about your system's processor(s).

Reference: <http://people.debian.org/~wouter/laptop/node21.html>

Incorrect Answers

- B: This file doesn't exist.
  - C: This file doesn't exist.
  - D: This file doesn't exist.
  - E: This file doesn't exist.
- 

**QUESTION 103**

The following shell script is run by cron on a regular basis:

```
x=$(find /home -name .rhost 2>/dev/null)
for i in $x; do
$(echo $I | cut -d/ -f3);z="$z $y"
rm $I;done
echo "Notice: $z" | mail root@example.com
```

Which best defines the action of this script?

- A. Verify the existence of users' .rhost files and removes the user.
- B. Remove all .rhosts files and notify each user of your action.
- C. Find all misplaced rhost files and remove them.
- D. This script checks for the existence of .rhost files, deletes them and reports the offending user names to root.
- E. Notify the root user of all .rhost files.

Answer: E