

LPI 117-201

Advanced Level Linux Certification
Version: 6.0

Topic 1, Volume A**QUESTION NO: 1**

What is the minimum number of disks required in a RAID5 array?

- A. 1
- B. 2
- C. 3
- D. 4
- E. 5

Answer: C

Explanation:

See <http://en.wikipedia.org/wiki/RAID>

RAID 5: Block-level striping with distributed parity. Minimum 3 disks.

QUESTION NO: 2

A server has three disks of 80GB each and must manage a database with 4 million records of 30KB each. The best configuration for this server, with the criteria being performance, configurability and flexibility, is:

- A. Use just LVM
- B. Use RAID1 over LVM
- C. Use LVM over RAID5
- D. Use LVM over RAID1
- E. Use RAID5 over LVM

Answer: C

Explanation:

RAID1 would not be performance enhancing (Mirroring), as would the default LVM (linear). So use a RAID5 for redundancy and performance and use LVM for added configurability and flexibility. A striped LVM only would work too, I guess, if you need no redundancy at all.

QUESTION NO: 3

The correct command sequence to create logical volumes on a Linux system is:

- A. lvcreate, pvcreate, vgcreate, mount, mke2fs
- B. pvcreate, vgcreate, lvcreate, mke2fs, mount
- C. vgcreate, lvcreate, pvcreate, mount, mke2fs
- D. mke2fs, pvcreate, vgcreate, lvcreate, mount
- E. pvcreate, lvcreate, vgcreate, mount, mke2fs

Answer: B

Explanation:

See <http://www.tldp.org/HOWTO/LVM-HOWTO/anatomy.html>

You need to create the Physical Volumes first, then the Volume Group (consisting of PVs), then the Logical Volumes, which you need to format and mount.

QUESTION NO: 4

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

- A. 18GB
- B. 36GB
- C. 54GB
- D. 72GB
- E. 90GB

Answer: C

Explanation:

See <http://en.wikipedia.org/wiki/RAID>

RAID5 capacity is available disks - 1. 5 disks - 1 - 1 spare = 3 disks. 3 x 18 = 54

QUESTION NO: 5

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. mklvm -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:

From the man pages:

lvcreate - create a logical volume in an existing volume group

-l, --extents LogicalExtentsNumber[%{VG|PVS|FREE}]

Gives the number of logical extents to allocate for the new logical volume. This can also be expressed as a percentage of the total space in the Volume Group with the suffix %VG, of the remaining free space in the Volume Group with the suffix %FREE, or of the remaining free space for the specified PhysicalVolume(s) with the suffix %PVS

-n, --name LogicalVolumeName The name for the new logical volume. Without this option a default names of "lvol#" will be generated where # is the LVM internal number of the logical volume.

QUESTION NO: 6

What is the purpose of vgextend?

- A. Extend a volume group's footprint on disk
- B. Add physical volumes to a volume group
- C. Increase the number of days between scheduled error checking
- D. Create a volume group which uses all available space on disk

Answer: B

Explanation:

From the man pages:

vgextend - add physical volumes to a volume group

QUESTION NO: 7

Which RAID level provides the most redundancy?

- A. RAID4
- B. RAID5
- C. RAID1
- D. RAID0

Answer: C

Explanation:

See <http://en.wikipedia.org/wiki/RAID>

- a,b) only 1 disk may fail, regardless how many disks are in the RAID
- d) no disk may fail
- c) n - 1 disks may fail

QUESTION NO: 8

Which pseudo-file contains information on the status of software RAID devices?

- A. /proc/raidstat
- B. /proc/mdstatus
- C. /proc/raidstatus
- D. /proc/mdstat
- E. /proc/raid/status

Answer: D

Explanation:

From the man pages:

/proc/mdstat If you're using the /proc filesystem, /proc/mdstat lists all active md devices with information about them. mdadm uses this to find arrays when --scan is given in Misc mode, and to monitor array reconstruction on Monitor mode.

QUESTION NO: 9

What information does the file modules.dep provide?

- A. A list of all modules available to the kernel.
- B. A list of modules required by each module.
- C. A list of devices and their module name.
- D. A list of modules compiled for this kernel.
- E. A list of modules the kernel needs to run.

Answer: B

Explanation:

From the man pages:

The modules.dep.bin as generated by module-init-tools depmod, lists the dependencies for every module in the directories under /lib/modules/version, where modules.dep.bin (or the human

readable version modules.

dep) is also located. It is used by utilities such as modprobe. The binary version will be used by default, if it was generated by a compatible version of depmod, with fallback to the generic modules.dep.

These files are not intended for editing or use by any additional utilities as their format is subject to change in the future. You should use the modinfo(8) command to obtain information about modules in a future proof and compatible fashion rather than touching these files.

QUESTION NO: 10

Before compiling a new kernel, what needs to be done?

- A. Compile kernel modules
- B. Configure the kernel options
- C. Delete old kernel sources
- D. Change to runlevel 1
- E. All of the above

Answer: B

Explanation:

See <http://www.faqs.org/docs/Linux-HOWTO/Kernel-HOWTO.html#impatient>

QUESTION NO: 11

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

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

Answer: D

Explanation:

/tmp and /var both need to be written to. /usr is the most static of the directories (excluding /usr/local)

QUESTION NO: 12

Which of the following best describes the format of the `/etc/fstab` file?

- A. mount point; device name; filesystem type; mount options; dump filesystem; fsck order
- B. device name; mount point; filesystem type; mount options; dump filesystem; fsck order
- C. mount point; filesystem type; device name; mount options; dump filesystem; fsck order
- D. device name; mount point; mount options; filesystem type; dump filesystem; fsck order
- E. device name; filesystem type; mount point; mount options; dump filesystem; fsck order

Answer: B

Explanation:

From the man pages:

The first field, (`fs_spec`), describes the block special device or remote filesystem to be mounted.

The second field, (`fs_file`), describes the mount point for the filesystem. For swap partitions, this field should be specified as `'none'`.

The third field, (`fs_vfstype`), describes the type of the filesystem.

The fourth field, (`fs_mntops`), describes the mount options associated with the filesystem.

The fifth field, (`fs_freq`), is used for these filesystems by the `dump(8)` command to determine which filesystems need to be dumped.

The sixth field, (`fs_passno`), is used by the `fsck(8)` program to determine the order in which filesystem checks are done at reboot time.

QUESTION NO: 13

Which of the following is a valid location in the automount configuration, when a Windows share needs to be made available to the system?

- A. `\\server\sharename`
- B. `://server/sharename`
- C. `server:/sharename`
- D. `:server/sharename`
- E. `\\\\server\\sharename`

Answer: B

Explanation:

From the man pages:

EXAMPLE

```
kernel -ro,soft,intr ftp.kernel.org:/pub/linux
boot -fstype=ext2 :/dev/hda1
windoze -fstype=smbfs ://windoze/c
removable -fstype=ext2 :/dev/hdd
cd -fstype=iso9660,ro :/dev/hdc
floppy -fstype=auto :/dev/fd0
server -rw,hard,intr / -ro myserver.me.org:/ \
/usr myserver.me.org:/usr \
/home myserver.me.org:/home
```

In the first line we have a NFS remote mount of the kernel directory on ftp.kernel.org. This is mounted readonly.

The second line mounts an ext2 volume on a local ide drive. The third makes a share exported from a

Windows machine available for automounting. The rest should be fairly self-explanatory. The last entry (thelast three lines) is an example of a multi-map (see below).

QUESTION NO: 14

What script is included with the kernel source to patch a kernel?

- A. patch
- B. patch-kernel
- C. apply-patch
- D. update-kernel

Answer: B

Explanation:

See <http://book.opensourceproject.org.cn/distrib/ubuntu/unleashed/opensource/0672329093/ch351ev1sec6.html>

It is possible to patch a kernel to the newest Linux kernel version as opposed to downloading the entire sourcecode. This choice can be beneficial for those who are not using a high-speed broadband connection. (A typical compressed kernel source file is nearly 30MB for a download time of about 10 minutes on a 512Kb DSL connection; adjust accordingly for your connection.) Whether you are patching existing sources or downloading the full source, the end results will be identical.

Patching the kernel is not a mindless task. It requires the user to retrieve all patches from her

current version to the version she wants to upgrade to. For example, if you are currently running 2.6.1 (and have those sources) and want to upgrade to 2.6.8, you must retrieve the 2.6.2 and 2.6.3 patch sets, and so on. Once downloaded, these patches must be applied in succession to upgrade to 2.6.8. This is more tedious than downloading the entire source, but useful for those who keep up with kernel hacking and want to perform incremental upgrades to keep their Linux kernel as up-to-date as possible.

To patch up to several versions in a single operation, you can use the patch-kernel script located in the kernel source directory for the kernel version you currently use. This script applies all necessary version patches to bring your kernel up to the latest version.

The format for using the patch-kernel script looks like this:

```
patch-kernel source_dir patch_dir stopversion
```

QUESTION NO: 15

Which command would you use to apply a diff file to an original?

- A. patch > diff-file
- B. patch < diff-file
- C. patch << diff-file
- D. cat diff-file >> kernel
- E. cat diff-file << kernel

Answer: B

Explanation:

From the man pages:

SYNOPSIS

```
patch [options] [originalfile [patchfile]]
```

but usually just

```
patch -pnum < patchfile
```

QUESTION NO: 16

One of the IDE hard drives in a server is transferring data very slowly. What command must be run to enable DMA on it?

- A. hdparm -d /dev/hda
- B. hdparm --dma /dev/hda
- C. hdparm --dma /dev/hda1
- D. hdparm -d /dev/hda1
- E. hdparm -d1 /dev/hda

Answer: E

Explanation:

hdparm provides a command line interface to various kernel interfaces supported by the Linux SATA/PATA/SAS "libata" subsystem and the older IDE driver subsystem. Many newer (2008 and later) USB drive enclosures now also support "SAT" (SCSI-ATA Command Translation) and therefore may also work with hdparm. Eg. recent WD "Passport" models and recent NexStar-3 enclosures. Some options may work correctly only with the latest kernels.

-d Disable/enable the "using_dma" flag for this drive. This option now works with most combinations of drives and PCI interfaces which support DMA and which are known to the kernel IDE driver. It is also a good idea to use the appropriate -X option in combination with -d1 to ensure that the drive itself is programmed for the correct DMA mode, although most BIOSs should do this for you at boot time. Using DMA nearly always gives the best performance, with fast I/O throughput and low CPU usage.

QUESTION NO: 17

What happens when the Linux kernel can't mount the root filesystem when booting?

- A. An error message is shown, showing which device couldn't be mounted or informing that init couldn't be found.
- B. An error message is shown and the system reboots after a keypress.
- C. An error message is shown and the system boots in maintenance mode.
- D. An error message is shown and the administrator is asked to specify a valid root filesystem to continue the boot process.
- E. An error message is shown, stating that the corresponding kernel module couldn't be loaded.

Answer: A

Explanation:

QUESTION NO: 18

You have written a little startup script (/usr/local/bin/startup.sh) that should automatically execute when entering runlevel 3. According to the standards, how can this best be achieved?