LPI Certification Self-Study Guide

David Horton

LPI Certification Self-Study Guide

David Horton

Abstract

The LPI Self-Study Guide is intended to provide a quick and inexpensive method for experienced Linux users to prepare for Linux Professional Institute (LPI) certification exams 101 and 102. The LPI Self-Study Guide is not a beginner's course and makes no attempt to teach any of the subject matter in detail. Instead the study guide provides a structured method for quickly reviewing the knowledge required by the exam objectives. Links to external references and documentation are provided for key terms and concepts and there are practice questions with answers at the end of each chapter.

Table of Contents

Introduction

1. About The LPI Self-Study Guide

There is no single, magical reference book that once read will enable a person to pass the Linux Professional Institute (LPI) certification exams. This LPI Self-Study Guide is no exception. The knowledge needed to pass the exam should come from a variety of sources including hands-on experience, instructor-led courses, man pages, HOWTO documents, books and self-study courses. This guide simply helps exam candidates structure their study efforts.

For long-time users of Linux systems this guide coupled with hands-on experience and man pages may be enough to prepare for the exams. Those who are new to Linux may be more comfortable using this guide as a final review after taking an instructor-led class.

2. About The Author

David Horton got started with GNU/Linux in 1996 when he needed a way to share a single dial-up Internet connection with his college room-mates. Dave is LPIC-1 certified and hold an M.S. degree in Information Systems from Roosevelt University, Chicago, Illinois USA as well as Cisco CCNA, CompTIA A+ and CompTIA Network+ certifications. Visit his web site at: http://www.happy-monkey.net.

3. Disclaimer

Using this document in no way guarantees that you will pass the LPI exams. This document is provided as-is with no warranty of any kind, either expressed or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Use the concepts, examples and information at your own risk. The author(s) do not take any responsibility for damages that may arise from the use of this document.

This document is not associated with nor endorsed by the Linux Professional Institute.

4. Copyright & License

This document is copyright (c) 2004-2006 by David Horton

This document is released under the terms of the Creative Commons Attribution-ShareAlike license [http://creativecommons.org/licenses/by-sa/2.0/]

5. Making Contributions

It is the author's wish that the practice questions be written and maintained by LPI certified persons. If you are LPI certified and would like to submit practice questions, please contact the author.

All Linux enthusiasts are welcome to submit references, point out errors and suggest improvements.

Any contributions submitted to the author will be covered under the same copyright and license as this document. By making contributions you signify that you agree to disclaim any copyright on the contribution and allow the contribution to be released under the terms of the Creative Commons Attribution-ShareAlike license [http://creativecommons.org/licenses/by-sa/2.0/].

Part I. General Information

Table of Contents

Chapter 1. Exam Preparation Methods

David Horton

1. Is The Self-Study Guide Right For You?

Self-study is not the easiest method of preparing for an exam. Look at the following checklist to help determine if the self-study guide is right for you.

- I am motivated, self-diciplined and good at managing my time.

- I prefer to learn things on my own rather than having someone teach me.

- I hold certifications from other vendors.

- I have used self-study to prepare for other certification tests and have been successful.

- I feel comfortable interacting with Linux using the command-line and often prefer it over a graphical user interface.

- I have used more than one Linux distribution.

- I am viewed as a Linux expert by my peers.

- I have visited LPI's web site [http://www.lpi.org/] and I feel comfortable with the exam objectives [http://www.lpi.org/en/lpic.html].

If you can answer yes to the majority of these questions self-study may be a good test preparation method for you. If you answered yes to less than half of the questions you may want to explore other, more formalized test preparation methods.

2. Alternatives To This Guide

If you are not confident in your ability to pass the LPI exams using this self-study guide you may want to find another test preparation option. There are many resources listed on LPI's web site [http://www.lpi.org/en/preparation.html].

Even if you are a self-study guru it is always a good idea to use a variety of sources for test preparation.

3. Study Groups

No matter if you choose self-study or formal classroom education it is smart to study in a group. Not only can you learn from others, but you can also reinforce your own knowledge by teaching. You may want to form a study group in your school or as part of your local Linux Users Group. There are also on-line study groups at the Lincert.com website [http://www.lincert.com/].

Chapter 2. Test Taking Tips

David Horton

1. Understanding the Test Format

When taking any exam it is important to understand the exam format prior to sitting down to take the test. This way you can be sure that you are being tested on your knowledge of the test material and not your understanding, or lack of understanding, of the test format.

The Linux Professional Institute (LPI) certification exams consist of fifty-one to seventy-three multiple choice and fill in the blank questions. The actual number of questions depends on which exam you are taking. LPI's exams are similar to computer-based exams from other vendors. If you have ever taken a Novell, Microsoft or Cisco certification test you will probably be very comfortable with the LPI exam format.

1.1. Types Of Questions You Will Encounter

LPI exam questions can be broken down into three basic types.

- Questions with fill in the blank answers

- Multiple choice questions with multiple correct answers

- Multiple choice questions with a single correct answer

All of the questions have a "mark for review option" that allows you to return to a question before finishing the exam. Mark for review is discussed in more detail in the next section.

Of the three types of questions, fill in the blank questions are probably the most challenging. There is no guessing, you must know the correct answer and enter it precisely. Be sure to read these questions very carefully and enter exactly what is asked for. Consider the following sample question:

```
System account information such as user ID and group ID is stored
in which file? (give the full path)
```

The correct answer is */etc/passwd*. An answer such as *passwd* is incorrect since the question clearly states that the full path should be given. This may seem like a trivial thing, but when you are under the stress of an exam situation it is tempting to rush to conclusions and not read the question fully. If you are not one-hundred percent sure that your answer is correct it is best to mark the question for review.

After dealing with fill in the blank questions you might think that multiple choice questions would be easy. None of the LPI questions is designed to be easy, but multiple choice does give you some advantage. Usually you can tackle a multiple choice question in small steps. Take a look at the following sample question:

```
Which of the following are valid IP addresses for use on the
Internet?  (choose two)

_ 192.16.15.211
_ 172.18.200.16
_ 68.143.255.10
_ 125.264.1.132
```

If you know the answer right away that is great, but if not there is a simple procedure you can follow to increase your chances of arriving at the correct answer.

The first thing you should do is to determine the type of multiple choice question. This sample question is a multiple answer type. You can figure this out by the fact that the question says, "choose two" and that the computer will allow you to select more than one answer.

Once the type of question is determined, the second step is to rule out any obviously wrong answers. In the case of the example above the fourth answer can be ruled out immediately, because 264 is above the valid range of values for an octet. (Octets are eight bits and can only be 0 - 255.)

At this point there are still three plausible looking answers and the question states that only two are correct. As a third step, re-read the question. Do not skim, but rather read slowly and deliberately looking for any information that may offer additional clues. The key phrase in the sample question is "for use on the Internet." This should tell you that the second answer is incorrect, because 172.18.200.16 belongs to the "class B reserved" range of addresses and cannot be used on the Internet.

The only remaining answers are the first and third choices. Since the question asks for two answers these should both be checked. If you are still unsure mark the question for review.

The other type of question you will encounter on the LPI exam is multiple choice single answer. These questions have one and only one correct answer. You can narrow down the field of potential correct answers using the same three steps as described for the multiple choice multiple answer questions, just know that in the end there can be only one correct answer. And if there is any doubt do not hesitate to mark the question for review.

1.2. Taking Advantage Of "Mark For Review"

In the previous section it was mentioned that questions should be marked for review if you are not one-hundred percent sure of the answer. This can be very beneficial. Marking questions allows you budget your time better by not spending too long on any one question. There is also the possibility that a later question will help you recall an elusive answer. This is particularly helpful for fill in the blank answers.

Once you have finished the last question on the exam you will be given a chance to revisit any marked questions. Take a look at the exam time clock and decide how much time you can afford to spend on each marked item. Trust your instincts when reviewing questions as many times your first choice is often the best choice. If an obvious answer is still not coming to you try to narrow down the selections if it is a multiple choice question. If all else fails take a guess. Guessing always gives better odds than simply leaving the question blank.

2. Managing Test Anxiety

No matter how long you have been using Linux or how much you study, you will be nervous on the day of the exam. Too much anxiety can work against you so it is important to take steps to minimize your stress level. There are several things you can do to calm yourself before the test.

2.1. In The Days Before The Exam

Before the exam day, try some of these tips.

- Study in a group if possible. Tutoring others is a great way to reinforce your own knowledge.

- Use more than one study aid. This guide is only one of many ways to prepare for the LPI exams. Additional resources can be found on the LPI web site [http://www.lpi.org].

- Consider taking the exam as part of an LPI exam lab. LPI offers exam sessions at certain Linux events at substantially discounted prices. Putting up $25 for an exam at an event is not nearly as stressful as gambling $100 at a testing center. Check the LPI web site [http://www.lpi.org] for details on exams lab events.

- Work with your body's natural clock. If you are a morning person schedule your exam early. If you are not a morning person schedule your exam in the afternoon.

- Make sure you know how to get to your testing location and how long it will take you to get there.

- Be sure you are properly prepared. If you feel you have not studied enough it may be possible to postpone the exam. Be sure to check with your testing center at least twenty-four hours in advance to see if this is an option.

2.2. On The Day Of The Exam

The following tips may help you manage your stress level on the day of the exam.

- Double-check the testing center rules and be sure to have proper ID with you.

- Make sure you are well rested and have had something to eat and drink before the test.

- If you are addicted to caffeine or nicotine be sure to get your fix before the exam.

- Arrive at the testing center early. Take time to decompress by walking the halls or visiting the water cooler. Take care of any bodily functions.

- Take time to adjust your chair and keyboard before clicking the exam start button.

- Try to relax by taking several deep breaths in through your nose and exhaling slowly and completely through your mouth.

- Visualize yourself passing the exam.

Part II. Exam 101

Table of Contents

Chapter 3. Hardware & Architecture

David Horton

1. A Brief Look At The Objectives

Successful completion of the Hardware & Architecture section of the LPI exam requires familiarity with Personal Computer (PC) [http://en.wikipedia.org/wiki/Personal_computer] expansion cards [http://en.wikipedia.org/wiki/Expansion_card] and peripherals [http://en.wikipedia.org/wiki/Peripherals] and how those devices interact with the system BIOS [http://en.wikipedia.org/wiki/BIOS] and the Linux Kernel [http://en.wikipedia.org/wiki/Linux_Kernel]. This includes understanding the difference between Industry Standard Architecture (ISA) [http://en.wikipedia.org/wiki/Industry_Standard_Architecture] and Peripheral Component Interconnect PCI [http://en.wikipedia.org/wiki/Peripheral_Component_Interconnect] as well as what interrupts (IRQs) [http://en.wikipedia.org/wiki/Interrupt], Input/Output (I/O) ports [http://en.wikipedia.org/wiki/Io_port] and Direct Memory Access (DMA) channels [http://en.wikipedia.org/wiki/Direct_memory_access] are used for. If you are the type of person whose PC spends as much time with the cover off as it does with the cover on you probably have a good head start on the preceding items. You will also need to know about Small Computer Systems Interface (SCSI) [http://en.wikipedia.org/wiki/SCSI] and Universal Serial Bus (USB) [http://en.wikipedia.org/wiki/USB] architectures, which are perhaps more difficult subjects. SCSI is more commonly used in high-end servers and may not be familiar to PC users. USB although more common in PC's is relatively new feature in the Linux kernel.

2. Ports, Interrupts and DMA

Expansion cards communicate with the PC using three basic methods, i/o ports, interrupts and DMA. The Linux kernel must keep track of these communication resources in order to properly interact with the expansion cards. The system administrator can see the kernel's view of i/o ports, interrupts and DMA channels by looking in the proc filesystem. There are three files in particular to note.

- `/proc/ioports`

- `/proc/interrupts`

- `/proc/dma`

All of these files can be viewed using the **cat** command. For example, **cat /proc/ioports** will show the the i/o ports used by the system. Take a moment to become familiar with the contents of these files. See if you can identify the various devices in your PC. Below are some excepts from the three files from a typical PC:

```
bash$ head -5 /proc/ioports
0000-001f : dma1
0020-003f : pic1
0040-005f : timer
0060-006f : keyboard
0070-007f : rtc

bash$ head -5 /proc/interrupts
           CPU0
  0:     204783          XT-PIC  timer
  1:       2410          XT-PIC  keyboard
  2:          0          XT-PIC  cascade
  5:       4004          XT-PIC  aic7xxx

bash$ cat /proc/dma
 4: cascade
```

Most modern expansion cards have their resources assigned automatically using plug-and-play (PnP) [http://en.wikipedia.org/wiki/Plug-n-play] technology and it is rare to find a card that requires manual configuration.

The way resources are assigned to a card depends upon which bus technology it uses. All PCI cards are designed to be plug-and-play and are assigned resources by the PC's BIOS during boot-up. The results of the configuration can be viewed using the **lspci** [http://unixhelp.ed.ac.uk/CGI/man-cgi?lspci+8] command. Older ISA cards may need to have their resources assigned by the operating system rather than the BIOS. Linux systems use the **isapnp** [http://www.roestock.demon.co.uk/isapnptools/isapnp.8.html] command and the corresponding `/etc/isapnp.conf` [http://www.roestock.demon.co.uk/isapnptools/isapnp.conf.5.html] file to set up ISA plug-and-play cards. The **pn-pdump** [http://www.roestock.demon.co.uk/isapnptools/pnpdump.8.html] command queries ISA plug-and-play cards to find their desired resource configuration and its output can be used to construct the `/etc/pnp.conf` file.

3. IDE Disks and LBA

Exploration of the `/proc/ide` directory reveals information about IDE devices present in the system. Of particular interest is the `/proc/ide/hda` directory since it contains information about the first, bootable IDE hard disk in the system. Two files in the `/proc/ide/hda` directory, `capacity` and `geometry`, are used to describe the size of the hard disk. The example below shows the contents of the files for an 80G hard disk.

```
bash# cat /proc/ide/hda/geometry
physical     38322/16/255
logical      9732/255/63

bash# cat /proc/ide/hda/capacity
156355584
```

The output of `geometry` shows the size in a cylinder/head/sector (CHS) format while `capacity` shows the size in logical block addressing (LBA) format. The older CHS format is limited to 1024 cylinders and can only describe disks up to 8G in size. Because of this limitation modern hard disks are almost always described using LBA. The Linux kernel uses LBA exclusively and provides the CHS parameters in `geometry` for informational purposes only. Because the operating system uses LBA it is important that the PC BIOS also be configured to use LBA.

4. SCSI Devices

In addition to IDE many high-performance systems use SCSI. Typical SCSI devices attached to a Linux system include hard drives, cd-roms and tape drives, but there can be others as well. All SCSI devices must be attached to a SCSI Host Adapter [http://en.wikipedia.org/wiki/Host_adapter] in order to interact with the system. The SCSI adapter is responsible for handling communication between the SCSI devices and the Linux kernel. Information about the host adapter and the devices attached to it will appear in the proc filesystem under the `/proc/scsi` directory. The following example shows a typical `/proc/scsi` directory.

```
bash# ls -F /proc/scsi
aic7xxx/   scsi
```

The directory `aic7xxx` contains information about the configuration of the host adapter (an Adaptec 2940 in this case.) There are many manufacturers of SCSI host adapters so the name and contents of the directory will vary depending on the partcular setup.

There is also a file named `/proc/scsi/scsi` that shows all SCSI devices as seen by the Linux kernel. An example is shown below.

```
bash# cat /proc/scsi/scsi
Attached devices:
Host: scsi0 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM      Model: DNES-309170W    Rev: SA30
  Type:   Direct-Access                   ANSI SCSI revision: 03
Host: scsi0 Channel: 00 Id: 04 Lun: 00
  Vendor: PHILIPS  Model: CDD2600         Rev: 1.07
  Type:   CD-ROM                          ANSI SCSI revision: 02
Host: scsi0 Channel: 00 Id: 05 Lun: 00
  Vendor: IOMEGA   Model: ZIP 100         Rev: E.03
  Type:   Direct-Access                   ANSI SCSI revision: 02
Host: scsi0 Channel: 00 Id: 06 Lun: 00
  Vendor: HP       Model: HP35480A        Rev: 1009
  Type:   Sequential-Access               ANSI SCSI revision: 02
```

Take a moment to look at the information contained in the file. Notice how each device on the SCSI bus has its own unique SCSI ID number. Hard disks are generally given lower SCSI ID's than CD-ROM's and tape drives and ID 0 is reserved for the bootable hard disk.

When accessing SCSI devices in Linux it is done using nodes in the `/dev` directory just like any other piece of hardware. Refering to the example above the IBM hard disk would be accessed as `/dev/sda` while the Iomega Zip disk is `/dev/sdb`. The Philips CD-ROM is `/dev/sr0` and the HP tape drive is `/dev/st0`.

5. USB Devices

Support for USB first appeared in Linux kernel version 2.2 and became much more robust in kernel version 2.4. USB is very similar to SCSI in many respects. The system has a USB host controller which functions much like a SCSI host adapter and USB storage devices appear as SCSI disks in the `/dev` directory.

In order for USB devices to be recognized there must be USB support in the kernel either compiled in or loaded as a module. In the case of a modular kernel the files required for basic USB support are as follows.

- `usbcore`

- `usb-uhci` or `usb-ohci` depending on the motherboard manufacturer.

The kernel will also need to have modules loaded for the particular USB device being used. For example `acm` for USB modems and `usb-storage` for storage devices like USB hard drives and USB CD-ROMs.

Many times it is also necessary to include hotplug support in the kernel because most USB devices are designed to be added and removed from the system without requiring a reboot. There is also a userspace program `/sbin/hotplug` [http://www.die.net/doc/linux/man/man8/hotplug.8.html] that helps the kernel deal with dynamically adding and removing USB devices. `/sbin/hotplug` does this by looking for a shell script (also called an agent) in the `/etc/hotplug` directory with the same name as the USB device being added or removed. The agent is responsible for handling the particulars of adding and removing the device.

6. Modems and Sound Cards

Modems and sound cards are given special attention on the LPI exam since they can be slightly more complicated than other types of hardware.

Modems have extra requirements to work properly with Linux and these are listed below.

- The modem must have the serial port set to the correct speed with setserial [http://unixhelp.ed.ac.uk/CGI/man-cgi?setserial+8].

- If the modem is a winmodem [http://en.wikipedia.org/wiki/Winmodem], the device must be supported.

In addition to setting up modem hardware the LPI exam also covers setting up a Point-to-Point Protocol (PPP) [http://en.wikipedia.org/wiki/Point-to-Point_Protocol] connection to an Internet provider. DSL and Cable modem users may want to refresh their memories by skimming the pppd [http://annys.eines.info/cgi-bin/man/man2html?8+pppd] and chat [http://annys.eines.info/cgi-bin/man/man2html?8+chat] man pages or the PPP-HOWTO [http://www.tldp.org/HOWTO/PPP-HOWTO/index.html].

Sound cards are typically difficult to configure because they use many different resources (i/o ports, interrupts & DMA channels) making it more likely that some sort of resource conflict will occur.

7. Practice Questions

You can gauge your familiarity with the subject of Hardware & Architecture by answering the practice questions below.

1. You are working with your hardware vendor's technical support people to troubleshoot a network card problem. The technical support representative wants to know which IRQ channel the card is using. Which file would tell you the IRQ channel for the card?

    A. /proc/interrupts

    B. /proc/ioports

    C. /proc/irqs

    D. /proc/sysconfig

2. You are providing telephone support to a novice user at a remote location, because the dial-in modem is not working. When you ask the user to tell you what port the modem is connected to he says, "COM1." What is the Linux device that corresponds to COM1? (provide the full path)

3. Management has finally approved the budget for you to buy your first high-capacity SCSI tape drive and SCSI host adapter. You install the host adapter in PCI slot 4 and set the tape drive's SCSI ID to 5. What Linux device would you use to access the new tape drive?

    A. /dev/st0

    B. /dev/st1

    C. /dev/st4

    D. /dev/st5

4. In which file would you look to find the base address of your system's sound card?

    A. /proc/sysconfig

    B. /proc/ioports

    C. /proc/interrupts

    D. /proc/base

5. What is the name of the file that contains information about direct memory access channels and the devices associated with them? (give the full path)

6. You have just purchased a 56kbps external modem and attached it to the PC serial port labled COM2. Communication with the modem is not working and /proc/ioports shows no serial devices. What might be the cause?

    A. The serial port is disabled in the BIOS

    B. The modem is turned off

    C. COM2 is reserved for serial mice

    D. The setserial command was not run properly at start-up

7. What BIOS feature can be turned on to enable the BIOS to access hard disks larger than 8 gigabytes?

    A. CHS

    B.   LBA

    C.   PIO

    D.   DMA

8.   Which of the following utilities can be used to create the configuration file `isapnp.conf`?

    A.   lspci

    B.   isapnp

    C.   lsmod

    D.   pnpdump

9.   In which kernel version were USB devices were first supported?

10. Which of the following kernel modules is required for USB? (choose 2)

    A.   hotplug

    B.   usbcore

    C.   usb-uhci

    D.   usbmgr

11. You have just purchased a second-hand PCI modem from an online auction. The modem does not appear to be working, but the seller assures you that it worked fine in his Windows PC. What could be the problem?

12. What is the name of the daemon that Linux systems use to establish a Point-to-Point connection? (specify the command name only with no options)

13. Which program can be used in conjuction with Point-to-Point connections to execute modem connection scripts non-interactively.

    A.   minicom

    B.   chat

    C.   pscript

    D.   uucico

8. Answers To Practice Questions

1.   The correct answer is A, `/proc/interrupts` shows IRQ lines and the devices using them. Answer B is incorrect because `/proc/ioports` will give the base address for the NIC card, not the IRQ. Answers C and D both refer to nonexistent files.

2.   The correct answer is `/dev/ttyS0`. Remember that Linux numbers devices starting from zero.

3.   The correct answer is A, `/dev/st0`. Answer B is incorrect because `/dev/st1` refers to the *second* SCSI tape drive and the question states that this is the *first* SCSI tape drive. Answers C and D are trying to trick you into thinking that the the device number is determined by the PCI slot number or the SCSI ID, but this is not the

case.

4. The correct answer is B, `/proc/ioports`. Answer C is incorrect because `/proc/interrupts` gives the IRQ channels used by the system. Answers A and D both refer to nonexistent files.

5. The correct answer is `/proc/dma`.

6. The correct answer is A, the modem is disabled in the BIOS. Answer B is incorrect because the modem is external and powering it off would not prevent the serial port device from appearing in `/proc/ioports`. Answer C is complete fiction as Linux does not reserve serial ports. Answer D is incorrect since **setserial** is used to set parameters such as handshaking and speed and not to enable or disable devices.

7. The correct answer is B, LBA. LBA stands for Logical Block Addressing and allows the BIOS to to access larger disks than could be accessed with the older CHS, or Cylinder Head Sector addressing. Answer A is incorrect for reasons just described. Answer C and D are incorrect because PIO and Ultra DMA both describe methods of data transfer, not data addressing.

8. Answer D is correct, **pnpdump** can be used to create the `isapnp.conf` file. Answer A, **lspci** gives information about devices on the PCI bus, not the ISA bus. The **isapnp** command cannot be used to generate its own configuration file so answer B is incorrect. Answer C, **lsmod**, is incorrect since this utility is used for listing kernel modules.

9. The correct answer is kernel 2.2. USB was not supported in kernels prior to this.

10. The correct answers are B and C. USB requires the modules `usbcore` and `usb-uhci`. Other motherboards might require `usb-ohci` in place of `usb-uhci` so do not be surprised if this is on the exam. Answers A and D are incorrect, because `hotplug` and `usbmgr` do not exist as modules. However, you will see `/etc/hotplug` and `/etc/usbmgr` directories which are used with USB.

11. Most likely the modem is a winmodem that is not supported by Linux.

12. The correct answer is **pppd**.

13. The correct answer is B, chat is used to run non-interactive login scripts for ppp connections. Answer A is incorrect, minicom is an *interactive* program. Answer C is a ficticious program. Answer D is incorrect, uucico is used for uucp not ppp.

9. Additional Resources

Listed below are some documents that may be helpful when preparing for the Hardware & Architecture portion of the LPI 101 exam.

•   The Linux kernel documentation contains a useful file called `devices.txt`. This file contains information about virtually every device one could expect to find in the `/dev` directory. Kernel documentation is generally found in the `/usr/src/linux/Documentation` directory.

•   The Large-Disk-HOWTO
    [http://www.ibiblio.org/pub/Linux/docs/HOWTO/other-formats/html_single/Large-Disk-HOWTO.html] does a good job of describing LBA, CHS and 1024 cylinder limitations.

•   The SCSI-2.4-HOWTO
    [http://www.ibiblio.org/pub/Linux/docs/HOWTO/other-formats/html_single/SCSI-2.4-HOWTO.html] goes into great detail about the Linux SCSI implentation.

•   The Linux USB web site maintains an FAQ [http://www.linux-usb.org/FAQ.html] that covers using USB devices with Linux.

- Modems and sound cards are covered in the older, but still relevant PPP-HOWTO
  [http://www.ibiblio.org/pub/Linux/docs/HOWTO/other-formats/html_single/PPP-HOWTO.html] and Sound-
  HOWTO [http://www.ibiblio.org/pub/Linux/docs/HOWTO/other-formats/html_single/Sound-HOWTO.html].

Chapter 4. Linux Installation & Package Management

1. A Brief Look At The Objectives

To be successful with the Linux Installation & Package Management topic of the exam, candidates should have a good understanding of disk partitioning [http://en.wikipedia.org/wiki/Partition_%28computing%29], installing a boot loader [http://en.wikipedia.org/wiki/Boot_loader] as well as the various methods of installing programs and shared libraries [http://en.wikipedia.org/wiki/Shared_libraries]. Your understanding of partitioning should go beyond just how to use fdisk [http://en.wikipedia.org/wiki/Fdisk] and focus on designing a partition scheme with separate /, /boot, /home, /usr, /var and swap file systems. Since LPI exams strive to be distribution neutral you should understand all of the various boot loader and package management options. This includes configuration of both LILO [http://en.wikipedia.org/wiki/LILO_%28boot_loader%29] and GRUB [http://en.wikipedia.org/wiki/GRand_Unified_Bootloader] boot loaders as well as installation of packages distributed in RPM [http://en.wikipedia.org/wiki/RPM_package_manager], DPKG [http://en.wikipedia.org/wiki/Dpkg], .tar.gz [http://en.wikipedia.org/wiki/Tar.gz] or source code [http://en.wikipedia.org/wiki/Source_code].

2. Disk Partitioning

Many times administrators of Linux systems will divide the hard disks into several partitions. It is common to see separate partitions for /, /home, /usr, /var and swap. Dividing a disk this way can help keep the root file system from filling up and crashing the system. For example, if /var/log/messages grows to an extremely large size only the partition holding /var will fill up.

If you have access to a system with multiple partitions, take a look at the sizes of various file systems by using the **df** command.

```
gnu-linux:~$ df
Filesystem          1K-blocks      Used Available Use% Mounted on
/dev/md0                50627     11995     36018  25% /
/dev/md1               199053    125479     63296  67% /usr
/dev/md2               249839     97970    138970  42% /var
/dev/md3               475953    299517    151848  67% /opt
/dev/md4              2111188    356140   1647804  18% /home
```

The actual sizes of the disk partitions can vary quite a bit from one system to another, but the sizes in relation to each other is fairly consistent. For example the partition holding / is small compared to the partitions for /usr, /var and /home. If you are inexperienced with this type of partitioning scheme or just want more information, refer to the Partition-HOWTO [http://www.tldp.org/HOWTO/Partition/].

Although not shown in the example above, it is sometimes necessary to create a separate /boot partition below the hard disk's 1024th cylinder. This is used as a work-around for older PC BIOSs that are unable to boot from a partition above the 1024th cylinder. Modern, LBA [http://en.wikipedia.org/wiki/Logical_block_addressing]-capable BIOSs do not suffer from this limitation. For additional information, see the Large-Disk-HOWTO [http://www.tldp.org/HOWTO/Large-Disk-HOWTO.html].

3. Boot Loaders

The two popular boot loaders for Linux systems are LILO and GRUB. All boot loaders for PC architecture work in essentially the same way.

1. The system powers up and the PC BIOS loads a 512-byte chunk of code called the Master Boot Record (MBR) from the first hard disk. This is the first stage.

2. Execution of the code from the MBR loads a larger, second-stage boot loader. This stage is not constrained to a 512-byte limit.

3. The second-stage boot loader then loads the operating system. For Linux systems this includes the Linux kernel and optionally an initial ramdisk (initrd) image.

3.1. GRUB

The GRUB boot loader stores copies of boot loader files in `/boot/grub`. Look at the example below and note the `stage1`, `stage2` and `menu.lst` files.

```
gnu-linux:/boot/grub$ ls -l
total 111
-rw-r--r--  1 root root    7840 Nov 17  2004 e2fs_stage1_5
-rw-r--r--  1 root root     480 Nov 17  2004 menu.lst
-rw-r--r--  1 root root     512 Nov 17  2004 stage1
-rw-r--r--  1 root root  101586 Nov 17  2004 stage2
```

When the system administrator runs **grub-install [http://www.die.net/doc/linux/man/man8/grub-install.8.html]** the `stage1` file is written to the first sector of the boot device. At boot-time the code from `stage1` will load and execute `stage2`. The configuration of kernel and ramdisk options is done by editing `menu.lst`.

Note

Some systems use a file called `grub.conf` in place of `menu.lst`. Do not be surprised to see these filenames used interchangeably in exam questions.

3.2. LILO

LILO loads the kernel and ramdisk in two stages similar to GRUB, but there are some important differences. LILO keeps some files in /boot, but its configuration file, `lilo.conf` [http://www.die.net/doc/linux/man/man5/lilo.conf.5.html], is stored in the `/etc` directory. LILO also requires administrators to run `/sbin/lilo` after any changes are made to `/etc/lilo.conf` or the kernel and ramdisk files.

3.3. MBR vs. Superblock

The first stage of the GRUB and LILO boot loaders does not always have to be installed in the Master Boot Record of the hard disk. Sometimes it is convenient to install the boot loader in the first sector, or superblock, of the root partition, particularly for systems that have multiple operating systems. Check the boot loader's man page for details on how to specify the installation location.

4. Package Management

In an effort to be distribution-neutral, the LPI exams cover several types of package management systems. This includes DPKG [http://en.wikipedia.org/wiki/Dpkg], RPM [http://en.wikipedia.org/wiki/RPM_Package_Manager] and .tar.gz [http://en.wikipedia.org/wiki/Tar.gz]. You should study the man pages for each and familiarize yourself with the command-line syntax required to perform the following tasks:

- Install a package

- Remove a package

- Upgrade a currently installed package with a new version

- Find out what files a particular package contains

- Given a filename, find out which package it came from

- Determine dependencies for a package

- Find the version number of a package

- Check a package's signature to verify its integrity.

The best way to learn package management for the exam is to do it on a day-to-day basis. Forget about the GUI tools that come with your system and start working exclusively on the command-line.

Unfortunately, you will probably only be able to get hands-on experience for two out of the three types of packages since DPKG and RPM are generally not seen together on the same system.

5. Shared Libraries

Shared libraries contain the code to perform common system tasks. Almost every program on a Linux system is linked to at least one shared libraries and cannot function independently of it. Most of the time everything works fine and the end user never even knows that shared libraries exist. However, a good system administrator, and a successful LPI candidate, must understand the inner workings of shared libraries.

The **ldd** command will let you find out which libraries a program uses. Take a look at the sample output below.

```
gnu-linux:~$ ldd /usr/bin/bash
        libncurses.so.5 => /usr/lib/libncurses.so.5 (0x40015000)
        libdl.so.2 => /lib/libdl.so.2 (0x40050000)
        libc.so.6 => /lib/libc.so.6 (0x40053000)
        /lib/ld-linux.so.2 => /lib/ld-linux.so.2 (0x40000000)
```

From the **ldd** output you can see that bash is linked to libncurses, libdl, libc and ld-linux. At the bottom of the list ld-linux, the dynamic loader, is in charge of locating and loading all of the other libraries. If ld-linux is unable to find any one of the other libraries in the list, bash will not run.

There are several methods ld-linux uses to find shared libraries. The easiest way to ensure a library can be found is to place it in either the /lib or /usr/lib directory since ld-linux will look here by default. Other directories to be searched can be added to the environmental variable LD_LIBRARY_PATH or the file /etc/ld.so.conf.

A portion of a typical ld.so.conf is shown below.

```
gnu-linux:~$ head -3 /etc/ld.so.conf
/usr/X11R6/lib/Xaw95
/usr/X11R6/lib/Xaw3d
/usr/X11R6/lib
```

New directories added to ld.so.conf do not take effect until the **ldconfig [http://www.die.net/doc/linux/man/man8/ldconfig.8.html]** command is executed and the changes are written to ld.so.cache. Most systems are configured to execute **ldconfig** at boot-up, but savvy Linux admins will execute it after making adjustments to ld.so.conf, because rebooting is for Windoze users.

6. Installing From Source Code

Occasionally, you may want to install a program that does not have a pre-compiled version available. When this happens you will need to know how to compile from source code. The basic steps involved are detailed below.

1. Expand the compressed archive using **bunzip2 [http://www.die.net/doc/linux/man/man1/bunzip2.1.html]** or **gunzip [http://www.die.net/doc/linux/man/man1/gunzip.1.html]**

2.   Extract the archive with **tar [http://www.die.net/doc/linux/man/man1/tar.1.html]**

3.   Set options by executing the `configure` script or manually editing the `Makefile`

4.   Run **make**

5.   Run **make install**

The best way to understand installing from source code is to do it. Download the hello program from ftp.gnu.org/gnu/hello [ftp://ftp.gnu.org/gnu/hello/] and practice the steps for installation.

7. Practice Questions
Note

Thanks go out to Paul van der Vlis for providing dpkg questions.

1.   You are planning a new Linux installation with separate partitions for /, /boot, /tmp and /usr. Which of the following file systems will be the largest?

   A.   /

   B.   /boot

   C.   /tmp

   D.   /usr

2.   What is the configuration file for LILO? (give the full path)

3.   What is the directory that contains configuration files for GRUB? (give the full path)

4.   You have accidentally deleted the file `/usr/lib/libm.so`. It needs to be re-installed, but you cannot remember what package it comes from. Which of the following commands would help you find the package that contains `libm.so`?

   A.   rpm -qf libm.so

   B.   rpm -e libm.so

   C.   rpm -ivh libm.so

   D.   rpm -qi libm.so

5.   You have accidentally deleted the file `/usr/lib/zlib.so` on a Debian system. Which of the following commands would help you find the package that contains `zlib.so`?

   A.   dpkg -S zlib.so

   B.   dpkg -r zlib.so

   C.   dpkg -i zlib.so

   D.   dpkg -l zlib.so

6.   Version 1.7 of your favorite web browser has just been released as an RPM package. You would like to install it while automatically un-installing any other versions. Which RPM command will allow you to install a new version of an RPM while automatically un-installing other versions?

  A. rpm --install

  B. rpm --upgrade

  C. rpm --verify

  D. rpm --erase

7. Version 1.5 of your favorite instant messenging client has just been released as a Debian package. You would like to install it while automatically un-installing any other versions. Which dpkg command will allow you to install a new version of a dpkg while automatically un-installing other versions?

  A. dpkg --extract

  B. dpkg --install

  C. dpkg --new

  D. dpkg --purge

8. What function will be performed by the command **rpm -ivh foo.rpm**?

  A. Verification of the files in `foo.rpm`

  B. Recalculation of the MD5 hash value for `foo.rpm`

  C. Installation of the package `foo.rpm`

  D. Verification of the signature for `foo.rpm`

9. What will happen if you issue the command **dpkg -iR /var/tmp/downloads**?

10. You have just downloaded the latest binary version of your favorite streaming audio server in a tarball called `llama-i386.tar.gz`. Which command could you use to extract the contents of `llama-i386.tar.gz`? (choose 2)

  A. tar -zxf llama-i386.tar.gz

  B. tar -xf llama-i386.tar.gz | gunzip -c

  C. gunzip -c llama-i386.tar.gz | tar xf -

  D. gunzip -c | tar xf - llama-i386.tar.gz

11. You have just installed the new wizbang-2.0 library and added its library path to `/etc/ld.so.conf`. What command should be run after adding the new library path to `ld.so.conf`?

8. Answers To Practice Questions

1. Of the partitions listed /usr will need to be the largest, so D is the correct answer. Answers A, B and C are incorrect, because the space requirements for / /boot and /tmp are small when compared to /usr.

2. LILO's configuration file is `/etc/lilo.conf`.

3. GRUB keeps configuration files in the `/boot/grub` directory.

4. The correct answer is A, **rpm -qf** will query a file to find out which package it came from. Answer B is incorrect because **rpm -e** is used to erase packages and cannot be used to find individual files. Answer C is incorrect since **rpm -ivh** is used to install packages, it has nothing to do with finding individual files. Answer D is incorrect because **rpm -qi** is used to query packages for information, not individual files.

5. The correct answer is A, **dpkg -S** will query a file to find out which package it came from. Answer B is incorrect because **dpkg -r** is used to remove packages and cannot be used to find individual files. Answer C is incorrect since **dpkg -i** is used to install packages, it has nothing to do with finding individual files. Answer D is incorrect because **dpkg -l** is used to list packages, not individual files.

6. The correct answer is B, **rpm --upgrade** will install the new version of an RPM and then un-install any other version. Answer A is incorrect, because while **rpm --install** will *install* a package it will not *un-install* other versions. Answer C is incorrect, because **rpm --verify** does not install or remove packages. Answer D is incorrect, **rpm --erase** removes packages.

7. The correct answer is B, **dpkg --install** will install the new version of a dpkg while removing any previous versions. Answer A is incorrect, because while **dpkg --extract** will *extract* a package it will not *un-install* other versions. Answer C is incorrect, because **dpkg --new** selects the binary format, but does not install anything. Answer D is incorrect, **dpkg --purge** removes packages.

8. The correct answer is C, **rpm -ivh foo.rpm** will install the package foo.rpm. More specifically it will install verbosely with hash marks to indicate the installation progress. Answer B is incorrect, do not confuse hash *marks* with hash *values*. Answer D is incorrect as signatures verification is performed with **rpm --checksig**.

9. The command **dpkg -iR ~/downloads** will recursively install all dpkg files found in the ~/downloads directory.

10. The correct answers are A and C. Both **tar -zxf llama-i386.tar.gz** and **gunzip -c llama-i386.tar.gz | tar xf -** will extract the contents of the llama-i386.tar.gz tarball. Answer B is incorrect, because it reverses the order of things by attempting to un-tar before decompressing. Answer D is incorrect, because the filename argument is in the wrong place and therefore **gunzip** will not pipe anything meaningful to **tar**.

11. The **ldconfig** command should be run after adding new library paths to /etc/ld.so.conf.

Chapter 5. GNU & Unix Commands

1. This Chapter Is Up For Adoption

If you feel you are an expert in the subject area for this LPI topic and would like to write this chapter please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide to communicate your intentions. Authors should be LPI certified or hold a similar certification in a related subject.

Please adopt only one chapter at a time and plan to finish the chapter within thirty days of adopting it.

More information about the study guide can be found at: http://www.happy-monkey.net/LPI/

1.1. Structuring The Chapter

The structure of chapters within the study guide needs to be consistent. The following structure may be used as a guide. Authors are also encouraged to view the chapter covering "Hardware & Architecture" as a template.

1. Author Name

    This is to ensure that authors get credit for their work. Please include your level of LPI certification after your name. If there are other appropriate certifications, please include these as well. For example, someone writing a chapter on networking might include the fact that they have a Cisco certification in addition to the LPI certification like this: "Joe Smith, LPIC-1, CCNA".

2. A Brief Look At The Objectives

    Each chapter should start with a high-level look at the LPI testing objectives. Within this "Brief Look" section there should be hyperlinks to sources of additional information. For example, one might say, "The networking section of the exam requires familiarity requires familiarity with IP addresses and subnet masks." The phrases 'IP addresses' and 'subnet masks' should be hyperlinked to additional information. Wikipedia is the prefered source of external information, but it is also helpful to link to man pages for specific commands.

3. Detailed Discussion

    Each one of the LPI exam objectives should receive a section for more detailed coverage of the objective and how it relates to Linux and the LPI exam. A good way to do this is by using hands-on exercises and examples from real-world Linux systems whenever possible. For example, when discussing a configuration file in /etc, give a sample of what a typical file might look like and encourage the reader to **cat** the file on their own system.

    Objectives may be grouped together when it is logical to do so.

4. Practice Questions

    At the end of each chapter there should be a section dedicated to practice questions that cover the exam objectives. *These should not be verbatim questions from the exam!* The idea is to let readers test their mastery of the knowledge in each chapter and get a feel for the exam format, not to help them cheat. Anyone intentionally submitting verbatim questions will be reported to the LPI and risks losing their certification.

1.2. DocBook XML Markup

The LPI Study Guide is written using DocBook XML markup. This allows the guide to be published in HTML, PDF and other formats using the same source document. Those familiar with DocBook are encouraged to submit chapters this way. Be sure to use markup for commands and filenames in addition to sections. Use the chapter on "Hardware

& Architecture" as a template.

For those not familiar with DocBook should submit chapters in plain text in a format similar to the example below.

```
A Brief Look At The Objectives

Successful completion of the Hardware & Architecture section
of the LPI exam requires familiarity with Personal Computer
(PC) expansion cards [http://url-to-wikipedia/expansion-cards]...

Next-Section-Title

More text...
```

1.3. Additional Questions

Please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide with any other questions not covered.

2. Practice Questions

1. Given the command **ls | tee dir**, where will the directory listing be sent?

    A.   To standard output and to standard error

    B.   To standard output and to a file called tee

    C.   To standard output and to a file called dir

    D.   To a file called tee and to a file called dir

2. A prospective employer is asking you to send a text file describing your qualifications in two-hundred words or less. What command will allow you to count the number of words in your text file?

3. Which command will add the directory /opt/bin to the end of your BASH shell's search path?

    A.   PATH=$PATH:/opt/bin

    B.   PATH=`cat $PATH /opt/bin`

    C.   PATH=/opt/bin

    D.   PATH="PATH:/opt/bin"

4. The command **unset FOO** does what?

    A.   Makes $FOO equal to zero

    B.   Makes $FOO equal to a null string

    C.   Removes any attributes assigned to FOO by **typeset**

    D.   Removes the shell variable FOO

5. Which of the following will send the standard error of **mkdir /tmp/foo** to `/dev/null`?

    A.   mkdir /tmp/foo </dev/null

    B.   mkdir /tmp/foo >/dev/null

  C. mkdir /tmp/foo 1>/dev/null

  D. mkdir /tmp/foo 2>/dev/null

6. Which of the following will display the string "oops" when the command **mkdir /tmp/bar** fails?

  A. mkdir /tmp/bar && echo "oops"

  B. mkdir /tmp/bar || echo "oops"

  C. mkdir /tmp/bar ; echo "oops"

  D. mkdir /tmp/bar | echo "oops"

7. What is the default signal for the **kill** command if no signal is specified as a command-line option?

  A. SIGHUP

  B. SIGINT

  C. SIGKILL

  D. SIGTERM

8. What signal can be used with the **kill** command to end processes that do not respond to the default signal?

  A. SIGHUP

  B. SIGINT

  C. SIGKILL

  D. SIGTERM

9. What command can be used to start a program with a lower scheduling priority?

10. Which of the following commands can be used to change the scheduling priority of a running process? (choose two)

  A. kill

  B. ps

  C. renice

  D. top

3. Answers To Practice Questions

1. The command **ls | tee dir** will send a directory listing to standard output and to the file called dir, so the correct answer is C. Answer A is wrong because the **tee** command does not send anything to standard error. Answer B is incorrect, because pipes send output to commands, not files. Answer D is incorrect, the proper way to send **ls** output to multiple files would be **ls | tee file1 file2**.

2. The **wc** command allows you to count the number of words in a text file.

3. The correct answer is A. Answer B is incorrect because **cat** is used to concatenate files, not variables and strings. Answer C is incorrect, it will ignore the previous PATH and set PATH to /opt/bin only. Answer D is incorrect because there is no $ in front of PATH to indicate it is a variable.

4. The correct answer is D, the command **unset FOO** removes the shell variable FOO. Answers A and B are incorrect since they refer to assigning values which is not what unset does. Answer C is incorrect since variable attributes are removed by using **typeset** with a plus instead of a minus in front of the attribute.

5. The correct answer is D, **command 2>/dev/null** redirects standard error to /dev/null. Remember the file descriptor for standard error is 2. Answer A is incorrect since it redirects standard input. Answers B and C both redirect standard output and are incorrect as well.

6. The correct answer is B, **command || echo "oops"** will display the string "oops" if command should fail. Answer A is incorrect since it will display "oops" when command is *successful*. Answer C is incorrect since it wil display "oops regardless of the success or failure of command. Answer D is incorrect, since a single "|" character is a pipeline not a compound command operator.

7. The default signal for the **kill** command is SIGTERM.

8. SIGKILL or -9 can be used when processes do not respond to the default SEGTERM.

9. The **nice** command can be used to start programs with lower scheduling priority. It can also be used to start programs with a higher priority, but only the superuser can do this.

10. The **renice** and **top** commands, answers C and D, can be used to change the priority of a running process. Answer A is incorrect, because **kill** is used to send signals to a process not change its scheduling priority. Answer B is incorrect, because **ps** only shows processes and cannot manipulate them.

4. References

The commands you can expect to see on the exam are listed on LPI's web site [http://www.lpi.org/] under Exam 101: Detailed Objectives [http://www.lpi.org/en/obj_101.html]. Be sure that you understand all of these commands.

4.1. Commands

Be sure that you are familiar with all of the commands listed for Topic 103 of the 101 exam objectives [http://www.lpi.org/en/obj_101.html]. Use the Linux manual pages [http://techpubs.sgi.com/tpl.cgi/linux/man/] and The GNU Linux Tools Summary [http://www.karakas-online.de/gnu-linux-tools-summary/] to help you learn about various commands you are unfamiliar with. Practice using the commands to perform various tasks on your system. If you are studying in a group make up sample tasks as challenges to the other group members.

4.2. Pipelines, Redirection and Compound Commands

See the BASH man page and info page for more information about how to redirect input, output and standard error as well as compound commands. Log into a Linux sytem and try entering commands like the answers shown above. Enter both correct and incorrect answers and view the results.

4.3. Processes and Priorities

See the man pages for ps(1), kill(1), nice(1), renice(1), signal(7) and top(1). Be sure you understand the difference between **nice** and **renice** and when it is appropriate to use one instead of the other. Pay attention to the priority adjustment numbers that can be used with **nice** and **renice** particularly what range of numbers is higher and which users are allowed to set higher priorities.

4.4. Foo and Bar

For more information on *foo* and *bar* refer to the wikipedia entry for foobar [http://en.wikipedia.org/wiki/Foobar]. To the author's knowledge the LPI exams do not feature any questions that use the terms *foo* or *bar*. However, many Linux references use these terms in examples and you may encounter them when studying.

Chapter 6. Devices, Linux Filesystems, Filesystem Hierarchy Standard

1. This Chapter Is Up For Adoption

If you feel you are an expert in the subject area for this LPI topic and would like to write this chapter please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide to communicate your intentions. Authors should be LPI certified or hold a similar certification in a related subject.

Please adopt only one chapter at a time and plan to finish the chapter within thirty days of adopting it.

More information about the study guide can be found at: http://www.happy-monkey.net/LPI/

1.1. Structuring The Chapter

The structure of chapters within the study guide needs to be consistent. The following structure may be used as a guide. Authors are also encouraged to view the chapter covering "Hardware & Architecture" as a template.

1. Author Name

    This is to ensure that authors get credit for their work. Please include your level of LPI certification after your name. If there are other appropriate certifications, please include these as well. For example, someone writing a chapter on networking might include the fact that they have a Cisco certification in addition to the LPI certification like this: "Joe Smith, LPIC-1, CCNA".

2. A Brief Look At The Objectives

    Each chapter should start with a high-level look at the LPI testing objectives. Within this "Brief Look" section there should be hyperlinks to sources of additional information. For example, one might say, "The networking section of the exam requires familiarity requires familiarity with IP addresses and subnet masks." The phrases 'IP addresses' and 'subnet masks' should be hyperlinked to additional information. Wikipedia is the prefered source of external information, but it is also helpful to link to man pages for specific commands.

3. Detailed Discussion

    Each one of the LPI exam objectives should receive a section for more detailed coverage of the objective and how it relates to Linux and the LPI exam. A good way to do this is by using hands-on exercises and examples from real-world Linux systems whenever possible. For example, when discussing a configuration file in `/etc`, give a sample of what a typical file might look like and encourage the reader to **cat** the file on their own system.

    Objectives may be grouped together when it is logical to do so.

4. Practice Questions

    At the end of each chapter there should be a section dedicated to practice questions that cover the exam objectives. *These should not be verbatim questions from the exam!* The idea is to let readers test their mastery of the knowledge in each chapter and get a feel for the exam format, not to help them cheat. Anyone intentionally submitting verbatim questions will be reported to the LPI and risks losing their certification.

1.2. DocBook XML Markup

The LPI Study Guide is written using DocBook XML markup. This allows the guide to be published in HTML, PDF and other formats using the same source document. Those familiar with DocBook are encouraged to submit chapters this way. Be sure to use markup for commands and filenames in addition to sections. Use the chapter on "Hardware & Architecture" as a template.

For those not familiar with DocBook should submit chapters in plain text in a format similar to the example below.

```
A Brief Look At The Objectives

Successful completion of the Hardware & Architecture section
of the LPI exam requires familiarity with Personal Computer
(PC) expansion cards [http://url-to-wikipedia/expansion-cards]...

Next-Section-Title

More text...
```

1.3. Additional Questions

Please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide with any other questions not covered.

2. Practice Questions

1. Which of the following commands can be used to create an ext2 (second extended) filesystem? (choose 2)

    A. ext2fs

    B. mke2fs

    C. mkfs.e2fs

    D. mkfs.ext2

2. The command **e2fsck** can be used to check which of the following types of filesystems? (choose 2)

    A. ext2

    B. ext3

    C. jfs

    D. vfat

3. Answers To Practice Questions

1. The correct answers are B and D, ext2 filesystems can be created with either **mke2fs** or **mkfs.ext2** commands. These are not actually two different programs, but rather **mkfs.ext2** is a link to **mke2fs**. Answers A and C are both incorrect as they refer to plausible looking but non-existent programs.

2. The correct answers are A and B, **e2fsck** can be used to check both ext2 (second extended) and ext3 (third extended) filesytems. Answers C and D are incorrect as jfs filesystems are checked with **jfs_fsck** and vfat filesystems are checked with **dosfsck**.

4. References

See the Description sections of the e2fsck(8) and mke2fs(8) man pages for more information about the similarities between ext2 and ext3 filesystems and the utilities used with them.

Chapter 7. The X Window System

David Horton

1. A Brief Look At The Objectives

Most Linux users take X-Windows [http://en.wikipedia.org/wiki/X-Windows] for granted; they turn on the machine and a Graphical User Interface (GUI) [http://en.wikipedia.org/wiki/Graphical_user_interface] magically appears. But, Linux system administrators, especially those studying for the LPI exams, must understand some of the inner workings of X-Windows. As an exam candidate you should have a general understanding of the system-wide X-Window system settings. In particular that these settings are kept in the /etc/X11/XF86Config [http://www.die.net/doc/linux/man/man5/xf86config.5.html] file and that you can use the xf86config [http://www.xfree86.org/current/XF86Config.5.html] utility (or the older XF86Setup [http://dell5.ma.utexas.edu/cgi-bin/man-cgi?XF86Setup+1]) to create an initial XF86Config. Another utility to re-member is xvidtune [http://www.die.net/doc/linux/man/man1/xvidtune.1.html] that allows fine tuning of the video modes defined in XF86Config. You also need to know a little about the functions performed by X Display Managers [http://en.wikipedia.org/wiki/X_display_manager] like XDM [http://en.wikipedia.org/wiki/XDM], KDM [http://en.wikipedia.org/wiki/KDM] and GDM [http://en.wikipedia.org/wiki/GNOME_Display_Manager] and where these display managers keep their configuration settings.

2. X-Windows Basics

X-Windows is the standard GUI for Linux distributions. (more to come)

3. X Display Managers

If you understand how `getty` [http://www.die.net/doc/linux/man/man8/agetty.8.html] and `login` [http://www.die.net/doc/linux/man/man1/login.1.html] work together to authenticate console-based users you should have no problem understanding X display managers. The display manager handles the tasks of connect-ing X-Windows to a screen, much like `getty` opens a TTY, and asking for user username and password information, much like `login` does when authenticating console-based users.

X display managers can also process network connections to the X server. Continuing with the console-based paral-lels you can thing of this functionality as being analogous to the telnet connection. On most machines remote con-nections are turned off for security reasons. This feature can be enabled by making a change in the display manager's configuration file and using the xhost [http://www.die.net/doc/linux/man/man1/xhost.1.html] utility to give the user and/or machine permission to connect.

The locations of the configuration files for the different display managers use a similar naming convention so they are easy to remember. The `/etc/X11/XDM` directory is used for the standard XDM while the KDE and Gnome dis-play managers use `/etc/X11/KDM` and `/etc/X11/GDM`, respectively. Just remember `/etc/X11/` followed by the three initials of the particular display manager.

4. Individual Customizations

Console-based programs often make use of hidden configuration files or environment variables [http://en.wikipedia.org/wiki/Environment_variable] that allow individual users to customize their experience with-out affecting other users on the system. For example, BASH reads settings from `.bashrc` in the users home direc-tory allowing users to override default settings for the shell. With so many parallels already seen between X-windows and the console-based login process is it any surprise that X uses similar method to override defaults?

There are a couple files in particular that are used by X-windows. There is `.xinitrc` that is similar in function to `.bashrc`. The `.xinitrc` file is read by the `xinit` program and can used to do things like start a particular program

when X starts up. You can think of this as a login script for X-Windows. There is also the `.XDefaults` file which allows individual users to customize the look and feel of many X applications. Some of the things that can be customized are positions of windows for various applications as well as window size fonts and colors used.

5. Practice Questions

under construction

6. Answers To Practice Questions

under construction

7. Additional References

The home page for X.Org [http://www.x.org/].

The home page for XFree86 [http://www.xfree86.org/].

The XDMCP-HOWTO [http://www.tldp.org/HOWTO/XDMCP-HOWTO/] from The Linux Documentation Project

Part III. Exam 102

Table of Contents

Chapter 8. Kernel

1. This Chapter Is Up For Adoption

If you feel you are an expert in the subject area for this LPI topic and would like to write this chapter please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide to communicate your intentions. Authors should be LPI certified or hold a similar certification in a related subject.

Please adopt only one chapter at a time and plan to finish the chapter within thirty days of adopting it.

More information about the study guide can be found at: http://www.happy-monkey.net/LPI/

1.1. Structuring The Chapter

The structure of chapters within the study guide needs to be consistent. The following structure may be used as a guide. Authors are also encouraged to view the chapter covering "Hardware & Architecture" as a template.

1. Author Name

   This is to ensure that authors get credit for their work. Please include your level of LPI certification after your name. If there are other appropriate certifications, please include these as well. For example, someone writing a chapter on networking might include the fact that they have a Cisco certification in addition to the LPI certification like this: "Joe Smith, LPIC-1, CCNA".

2. A Brief Look At The Objectives

   Each chapter should start with a high-level look at the LPI testing objectives. Within this "Brief Look" section there should be hyperlinks to sources of additional information. For example, one might say, "The networking section of the exam requires familiarity requires familiarity with IP addresses and subnet masks." The phrases 'IP addresses' and 'subnet masks' should be hyperlinked to additional information. Wikipedia is the prefered source of external information, but it is also helpful to link to man pages for specific commands.

3. Detailed Discussion

   Each one of the LPI exam objectives should receive a section for more detailed coverage of the objective and how it relates to Linux and the LPI exam. A good way to do this is by using hands-on exercises and examples from real-world Linux systems whenever possible. For example, when discussing a configuration file in /etc, give a sample of what a typical file might look like and encourage the reader to **cat** the file on their own system.

   Objectives may be grouped together when it is logical to do so.

4. Practice Questions

   At the end of each chapter there should be a section dedicated to practice questions that cover the exam objectives. *These should not be verbatim questions from the exam!* The idea is to let readers test their mastery of the knowledge in each chapter and get a feel for the exam format, not to help them cheat. Anyone intentionally submitting verbatim questions will be reported to the LPI and risks losing their certification.

1.2. DocBook XML Markup

The LPI Study Guide is written using DocBook XML markup. This allows the guide to be published in HTML, PDF and other formats using the same source document. Those familiar with DocBook are encouraged to submit chapters this way. Be sure to use markup for commands and filenames in addition to sections. Use the chapter on "Hardware

& Architecture" as a template.

For those not familiar with DocBook should submit chapters in plain text in a format similar to the example below.

```
A Brief Look At The Objectives

Successful completion of the Hardware & Architecture section
of the LPI exam requires familiarity with Personal Computer
(PC) expansion cards [http://url-to-wikipedia/expansion-cards]...

Next-Section-Title

More text...
```

1.3. Additional Questions

Please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide with any other questions not covered.

2. Practice Questions

1. The system is experiencing network connectivity problems and you have been called in to troubleshoot the situation. You suspect that the kernel module for the ethernet card did not load properly during start-up. What command would you use to find out if the module is currently loaded?

2. You have just downloaded the source code for kernel 2.4.26 and unpacked it into /usr/src. After changing directory to /usr/src/linux-2.4.26, what is the next command you might run.

    A. make install

    B. ./configure --prefix=/boot

    C. make menuconfig

    D. lilo

3. Answers To Practice Questions

1. **lsmod** would be the correct command to show modules that are currently loaded.

2. Answer C, **make menuconfig**, is the correct answer. Answer A is incorrect since the kernel's Makefile does not have an install target. Answer B is incorrect since the kernel does not use the autoconfig system. Answer D installs the LILO boot loader, it does not do anything to the kernel.

Chapter 9. Boot, Initialization, Shutdown and Runlevels

1. This Chapter Is Up For Adoption

If you feel you are an expert in the subject area for this LPI topic and would like to write this chapter please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide to communicate your intentions. Authors should be LPI certified or hold a similar certification in a related subject.

Please adopt only one chapter at a time and plan to finish the chapter within thirty days of adopting it.

More information about the study guide can be found at: http://www.happy-monkey.net/LPI/

1.1. Structuring The Chapter

The structure of chapters within the study guide needs to be consistent. The following structure may be used as a guide. Authors are also encouraged to view the chapter covering "Hardware & Architecture" as a template.

1. Author Name

   This is to ensure that authors get credit for their work. Please include your level of LPI certification after your name. If there are other appropriate certifications, please include these as well. For example, someone writing a chapter on networking might include the fact that they have a Cisco certification in addition to the LPI certification like this: "Joe Smith, LPIC-1, CCNA".

2. A Brief Look At The Objectives

   Each chapter should start with a high-level look at the LPI testing objectives. Within this "Brief Look" section there should be hyperlinks to sources of additional information. For example, one might say, "The networking section of the exam requires familiarity requires familiarity with IP addresses and subnet masks." The phrases 'IP addresses' and 'subnet masks' should be hyperlinked to additional information. Wikipedia is the prefered source of external information, but it is also helpful to link to man pages for specific commands.

3. Detailed Discussion

   Each one of the LPI exam objectives should receive a section for more detailed coverage of the objective and how it relates to Linux and the LPI exam. A good way to do this is by using hands-on exercises and examples from real-world Linux systems whenever possible. For example, when discussing a configuration file in `/etc`, give a sample of what a typical file might look like and encourage the reader to **cat** the file on their own system.

   Objectives may be grouped together when it is logical to do so.

4. Practice Questions

   At the end of each chapter there should be a section dedicated to practice questions that cover the exam objectives. *These should not be verbatim questions from the exam!* The idea is to let readers test their mastery of the knowledge in each chapter and get a feel for the exam format, not to help them cheat. Anyone intentionally submitting verbatim questions will be reported to the LPI and risks losing their certification.

1.2. DocBook XML Markup

The LPI Study Guide is written using DocBook XML markup. This allows the guide to be published in HTML, PDF and other formats using the same source document. Those familiar with DocBook are encouraged to submit chapters this way. Be sure to use markup for commands and filenames in addition to sections. Use the chapter on "Hardware & Architecture" as a template.

For those not familiar with DocBook should submit chapters in plain text in a format similar to the example below.

```
A Brief Look At The Objectives

Successful completion of the Hardware & Architecture section
of the LPI exam requires familiarity with Personal Computer
(PC) expansion cards [http://url-to-wikipedia/expansion-cards]...

Next-Section-Title

More text...
```

1.3. Additional Questions

Please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide with any other questions not covered.

2. Practice Questions

1. Your /etc/inittab file has a line that reads "id:1:initdefault:" What mode will the system enter when it is booted?

    A. Single user mode

    B. Non-network multi-user mode

    C. Network multi-user mode

    D. Multi-user GUI mode

2. According to the Linux Standard Base (LSB) specification, runlevel 3 is reserved for which of the following system states?

    A. Single user mode

    B. Non-network multi-user mode

    C. Network multi-user mode

    D. Multi-user GUI mode

3. You have edited your /etc/inittab and changed the line "id:5:initdefault:" to read "su:5:initdefault". What will be the runlevel on the next reboot?

    A. 0

    B. 1

    C. 5

    D. None, the init daemon will prompt for the runlevel before finishing boot-up.

4. What command can be used to replay messages that were displayed while the Linux kernel was kernel booting?

3. Answers to Practice Questions

1. Answer A is correct since runlevel 1 is single-user mode. Answer B, C and D are incorrect. According to the Linux Standard Base specification, non-networked multi-user mode is runlevel 2, networked multi-user mode is runlevel 3 and multi-user GUI mode is runlevel 5.

2. The correct answer is C, runlevel 3 is network multi-user mode on an LSB compliant system. Single user mode is runlevel 1, non-network multi-user mode is runlevel 2 and multi-user GUI mode is runlevel 5.

3. Runlevel 5, answer C, is correct. Answer A, runlevel 0, is for system halt. Although answer B may be tempting it is incorrect. The label of "su" makes no difference in the runlevel, everything is determined by the number "5". Answer D is incorrect since the system will only prompt for a runlevel if initdefault is missing.

4. The **dmesg** command will allow you to print the kernel ring-buffer which contains the kernel's boot messages.

4. References

For an in-depth reference see the From PowerUp To BASH Prompt HOWTO
[http://www.tldp.org/HOWTO/From-PowerUp-To-Bash-Prompt-HOWTO.html] from The Linux Documentation
Project [http://www.tldp.org/]. Also skim the run levels section of the linux standard base
[http://www.linuxbase.org/spec/] specification document.

Chapter 10. Printing

1. This Chapter Is Up For Adoption

If you feel you are an expert in the subject area for this LPI topic and would like to write this chapter please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide to communicate your intentions. Authors should be LPI certified or hold a similar certification in a related subject.

Please adopt only one chapter at a time and plan to finish the chapter within thirty days of adopting it.

More information about the study guide can be found at: http://www.happy-monkey.net/LPI/

1.1. Structuring The Chapter

The structure of chapters within the study guide needs to be consistent. The following structure may be used as a guide. Authors are also encouraged to view the chapter covering "Hardware & Architecture" as a template.

1. Author Name

   This is to ensure that authors get credit for their work. Please include your level of LPI certification after your name. If there are other appropriate certifications, please include these as well. For example, someone writing a chapter on networking might include the fact that they have a Cisco certification in addition to the LPI certification like this: "Joe Smith, LPIC-1, CCNA".

2. A Brief Look At The Objectives

   Each chapter should start with a high-level look at the LPI testing objectives. Within this "Brief Look" section there should be hyperlinks to sources of additional information. For example, one might say, "The networking section of the exam requires familiarity requires familiarity with IP addresses and subnet masks." The phrases 'IP addresses' and 'subnet masks' should be hyperlinked to additional information. Wikipedia is the prefered source of external information, but it is also helpful to link to man pages for specific commands.

3. Detailed Discussion

   Each one of the LPI exam objectives should receive a section for more detailed coverage of the objective and how it relates to Linux and the LPI exam. A good way to do this is by using hands-on exercises and examples from real-world Linux systems whenever possible. For example, when discussing a configuration file in /etc, give a sample of what a typical file might look like and encourage the reader to **cat** the file on their own system.

   Objectives may be grouped together when it is logical to do so.

4. Practice Questions

   At the end of each chapter there should be a section dedicated to practice questions that cover the exam objectives. *These should not be verbatim questions from the exam!* The idea is to let readers test their mastery of the knowledge in each chapter and get a feel for the exam format, not to help them cheat. Anyone intentionally submitting verbatim questions will be reported to the LPI and risks losing their certification.

1.2. DocBook XML Markup

The LPI Study Guide is written using DocBook XML markup. This allows the guide to be published in HTML, PDF and other formats using the same source document. Those familiar with DocBook are encouraged to submit chapters this way. Be sure to use markup for commands and filenames in addition to sections. Use the chapter on "Hardware

& Architecture" as a template.

For those not familiar with DocBook should submit chapters in plain text in a format similar to the example below.

```
A Brief Look At The Objectives

Successful completion of the Hardware & Architecture section
of the LPI exam requires familiarity with Personal Computer
(PC) expansion cards [http://url-to-wikipedia/expansion-cards]...

Next-Section-Title

More text...
```

1.3. Additional Questions

Please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide with any other questions not covered.

2. Practice Questions

1. Which of the following commands can be used to show the status of a print queue? (choose two)

    A. lpr

    B. lpq

    C. lpc

    D. lpd

2. Which command can be used to remove a print job from a queue?

3. It is time to close down the computer lab and you would like to stop anyone from printing any last minute reports. Which command allows you to disable all printers?

3. Answers To Practice Questions

1. The correct answers are B and C. Both **lpq** and **lpc** can be used to show the status of print queues. Answer A is incorrect, because **lpr** is only used to submit jobs. Answer D is incorrect, **lpd** is the line printer daemon that runs in the background.

2. The correct answer is **lprm**.

3. The correct answer is **lpc**.

4. References

For information of printing see the man pages for **lpc**, **lpd**, **lpq**, **lpr** and **lprm**. For a more in-depth discussion of Linux printing see Patrick Powell's LPRng reference manual available from the LPRng home page [http://www.lprng.org/].

Chapter 11. Documentation

1. This Chapter Is Up For Adoption

If you feel you are an expert in the subject area for this LPI topic and would like to write this chapter please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide to communicate your intentions. Authors should be LPI certified or hold a similar certification in a related subject.

Please adopt only one chapter at a time and plan to finish the chapter within thirty days of adopting it.

More information about the study guide can be found at: http://www.happy-monkey.net/LPI/

1.1. Structuring The Chapter

The structure of chapters within the study guide needs to be consistent. The following structure may be used as a guide. Authors are also encouraged to view the chapter covering "Hardware & Architecture" as a template.

1. Author Name

   This is to ensure that authors get credit for their work. Please include your level of LPI certification after your name. If there are other appropriate certifications, please include these as well. For example, someone writing a chapter on networking might include the fact that they have a Cisco certification in addition to the LPI certification like this: "Joe Smith, LPIC-1, CCNA".

2. A Brief Look At The Objectives

   Each chapter should start with a high-level look at the LPI testing objectives. Within this "Brief Look" section there should be hyperlinks to sources of additional information. For example, one might say, "The networking section of the exam requires familiarity requires familiarity with IP addresses and subnet masks." The phrases 'IP addresses' and 'subnet masks' should be hyperlinked to additional information. Wikipedia is the prefered source of external information, but it is also helpful to link to man pages for specific commands.

3. Detailed Discussion

   Each one of the LPI exam objectives should receive a section for more detailed coverage of the objective and how it relates to Linux and the LPI exam. A good way to do this is by using hands-on exercises and examples from real-world Linux systems whenever possible. For example, when discussing a configuration file in /etc, give a sample of what a typical file might look like and encourage the reader to **cat** the file on their own system.

   Objectives may be grouped together when it is logical to do so.

4. Practice Questions

   At the end of each chapter there should be a section dedicated to practice questions that cover the exam objectives. *These should not be verbatim questions from the exam!* The idea is to let readers test their mastery of the knowledge in each chapter and get a feel for the exam format, not to help them cheat. Anyone intentionally submitting verbatim questions will be reported to the LPI and risks losing their certification.

1.2. DocBook XML Markup

The LPI Study Guide is written using DocBook XML markup. This allows the guide to be published in HTML, PDF and other formats using the same source document. Those familiar with DocBook are encouraged to submit chapters this way. Be sure to use markup for commands and filenames in addition to sections. Use the chapter on "Hardware

& Architecture" as a template.

For those not familiar with DocBook should submit chapters in plain text in a format similar to the example below.

```
A Brief Look At The Objectives

Successful completion of the Hardware & Architecture section
of the LPI exam requires familiarity with Personal Computer
(PC) expansion cards [http://url-to-wikipedia/expansion-cards]...

Next-Section-Title

More text...
```

1.3. Additional Questions

Please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide with any other questions not covered.

2. Practice Questions

1. The command **man -k passwd** gives the same results as which one of the following commands:

    A. whatis passwd

    B. apropos passwd

    C. passwd --help

    D. info passwd

3. Answers To Practice Questions

1. The correct answer is B, **man -k** and **apropos** are equivalent. Answer A is incorrect, **whatis** is the same as **man -f**. Answer C is incorrect, **passwd --help** gives a brief listing of command-line options and is unrelated to man pages. Answer D is incorrect, info pages are part of a different documentation tool.

4. References

See the manual page for the **man** command.

Chapter 12. Shells, Scripting, Programming and Compiling

1. This Chapter Is Up For Adoption

If you feel you are an expert in the subject area for this LPI topic and would like to write this chapter please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide to communicate your intentions. Authors should be LPI certified or hold a similar certification in a related subject.

Please adopt only one chapter at a time and plan to finish the chapter within thirty days of adopting it.

More information about the study guide can be found at: http://www.happy-monkey.net/LPI/

1.1. Structuring The Chapter

The structure of chapters within the study guide needs to be consistent. The following structure may be used as a guide. Authors are also encouraged to view the chapter covering "Hardware & Architecture" as a template.

1. Author Name

   This is to ensure that authors get credit for their work. Please include your level of LPI certification after your name. If there are other appropriate certifications, please include these as well. For example, someone writing a chapter on networking might include the fact that they have a Cisco certification in addition to the LPI certification like this: "Joe Smith, LPIC-1, CCNA".

2. A Brief Look At The Objectives

   Each chapter should start with a high-level look at the LPI testing objectives. Within this "Brief Look" section there should be hyperlinks to sources of additional information. For example, one might say, "The networking section of the exam requires familiarity requires familiarity with IP addresses and subnet masks." The phrases 'IP addresses' and 'subnet masks' should be hyperlinked to additional information. Wikipedia is the prefered source of external information, but it is also helpful to link to man pages for specific commands.

3. Detailed Discussion

   Each one of the LPI exam objectives should receive a section for more detailed coverage of the objective and how it relates to Linux and the LPI exam. A good way to do this is by using hands-on exercises and examples from real-world Linux systems whenever possible. For example, when discussing a configuration file in `/etc`, give a sample of what a typical file might look like and encourage the reader to **cat** the file on their own system.

   Objectives may be grouped together when it is logical to do so.

4. Practice Questions

   At the end of each chapter there should be a section dedicated to practice questions that cover the exam objectives. *These should not be verbatim questions from the exam!* The idea is to let readers test their mastery of the knowledge in each chapter and get a feel for the exam format, not to help them cheat. Anyone intentionally submitting verbatim questions will be reported to the LPI and risks losing their certification.

1.2. DocBook XML Markup

The LPI Study Guide is written using DocBook XML markup. This allows the guide to be published in HTML, PDF and other formats using the same source document. Those familiar with DocBook are encouraged to submit chapters this way. Be sure to use markup for commands and filenames in addition to sections. Use the chapter on "Hardware & Architecture" as a template.

For those not familiar with DocBook should submit chapters in plain text in a format similar to the example below.

```
A Brief Look At The Objectives

Successful completion of the Hardware & Architecture section
of the LPI exam requires familiarity with Personal Computer
(PC) expansion cards [http://url-to-wikipedia/expansion-cards]...

Next-Section-Title

More text...
```

1.3. Additional Questions

Please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide with any other questions not covered.

Chapter 13. Administrative Tasks

1. This Chapter Is Up For Adoption

If you feel you are an expert in the subject area for this LPI topic and would like to write this chapter please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide to communicate your intentions. Authors should be LPI certified or hold a similar certification in a related subject.

Please adopt only one chapter at a time and plan to finish the chapter within thirty days of adopting it.

More information about the study guide can be found at: http://www.happy-monkey.net/LPI/

1.1. Structuring The Chapter

The structure of chapters within the study guide needs to be consistent. The following structure may be used as a guide. Authors are also encouraged to view the chapter covering "Hardware & Architecture" as a template.

1. Author Name

    This is to ensure that authors get credit for their work. Please include your level of LPI certification after your name. If there are other appropriate certifications, please include these as well. For example, someone writing a chapter on networking might include the fact that they have a Cisco certification in addition to the LPI certification like this: "Joe Smith, LPIC-1, CCNA".

2. A Brief Look At The Objectives

    Each chapter should start with a high-level look at the LPI testing objectives. Within this "Brief Look" section there should be hyperlinks to sources of additional information. For example, one might say, "The networking section of the exam requires familiarity requires familiarity with IP addresses and subnet masks." The phrases 'IP addresses' and 'subnet masks' should be hyperlinked to additional information. Wikipedia is the prefered source of external information, but it is also helpful to link to man pages for specific commands.

3. Detailed Discussion

    Each one of the LPI exam objectives should receive a section for more detailed coverage of the objective and how it relates to Linux and the LPI exam. A good way to do this is by using hands-on exercises and examples from real-world Linux systems whenever possible. For example, when discussing a configuration file in /etc, give a sample of what a typical file might look like and encourage the reader to **cat** the file on their own system.

    Objectives may be grouped together when it is logical to do so.

4. Practice Questions

    At the end of each chapter there should be a section dedicated to practice questions that cover the exam objectives. *These should not be verbatim questions from the exam!* The idea is to let readers test their mastery of the knowledge in each chapter and get a feel for the exam format, not to help them cheat. Anyone intentionally submitting verbatim questions will be reported to the LPI and risks losing their certification.

1.2. DocBook XML Markup

The LPI Study Guide is written using DocBook XML markup. This allows the guide to be published in HTML, PDF and other formats using the same source document. Those familiar with DocBook are encouraged to submit chapters this way. Be sure to use markup for commands and filenames in addition to sections. Use the chapter on "Hardware

& Architecture" as a template.

For those not familiar with DocBook should submit chapters in plain text in a format similar to the example below.

```
A Brief Look At The Objectives

Successful completion of the Hardware & Architecture section
of the LPI exam requires familiarity with Personal Computer
(PC) expansion cards [http://url-to-wikipedia/expansion-cards]...

Next-Section-Title

More text...
```

1.3. Additional Questions

Please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide with any other questions not covered.

Chapter 14. Networking Fundamentals

David Horton

1. A Brief Look At The Objectives

To be successful with the network fundamentals section of the LPI exam candidates should possess a good working knowledge of Internet Protocol (IP) [http://en.wikipedia.org/wiki/Internet_Protocol] (version 4) networking. This includes understanding the concepts of address classes [http://en.wikipedia.org/wiki/Classful_network], subnets [http://en.wikipedia.org/wiki/Subnetwork] and private IP address [http://en.wikipedia.org/wiki/Private_IP_address] ranges. It is also necessary to know port numbers [http://en.wikipedia.org/wiki/List_of_well-known_ports_%28computing%29] for popular network services and common protocols [http://en.wikipedia.org/wiki/Internet_protocol_suite] like Transmission Control Protocol (TCP) [http://en.wikipedia.org/wiki/Transmission_Control_Protocol], User Datagram Protocol (UDP) [http://en.wikipedia.org/wiki/User_Datagram_Protocol] and Internet Control Message Protocol (ICMP) [http://en.wikipedia.org/wiki/Internet_Control_Message_Protocol]. This section of the exam will also cover commands and configuration files associated with networking as well as utilities for troubleshooting connectivity.

2. IP Basics

under construction

3. IP Ports and Protocols

Making connections to internet services requires knowing more than just an IP address. There is also a proper port number and protocol that must be used. Most end-users are blissfully unaware of this fact because the operating system takes care of the mundane details.

For Linux systems the `/etc/services` file is responsible for translating a service name, like FTP or telnet, into port numbers, like 21 or 23. An excerpt of `/etc/services` is show below.

```
bash$ sed -n 43,51p /etc/services
ftp             21/tcp
ftp             21/udp          fsp fspd
ssh             22/tcp                          # SSH Remote Login Protocol
ssh             22/udp                          # SSH Remote Login Protocol
telnet          23/tcp
telnet          23/udp
# 24 - private mail system
smtp            25/tcp          mail
smtp            25/udp          mail
```

Some of the well-known port numbers that may appear on the exam are listed below. Take time to memorize this entire list.

- 20 - File Transfer Protocol (FTP) data

- 21 - File Transfer Protocol (FTP) control

- 23 - Telnet

- 25 - Simple Mail Transfer Protocol (SMTP)

- 53 - Domain Name Service (DNS)

- 80 - HyperText Transfer Protocol (HTTP)

- 110 - Post Office Protocol version 3 (POP3)

- 119 - Network News Transport Protocol (NNTP)

- 139 - NetBIOS Session Service

- 143 - Internet Message Access Protocol (IMAP)

- 161 - Simple Network Management Protocol (SNMP)

In addition to the ip address and port number that are needed to make a connection to an internet service it is also necessary to use the correct protocol. This is another detail that is transparent to the end user, because it is handled by the operating system. Information about the protocols used with IP are kept in the `/etc/protocols` file. An excerpt is shown below.

```
bash$ sed -n -e 12,19p -e 30p /etc/protocols
ip      0       IP            # internet protocol, pseudo protocol number
#hopopt 0       HOPOPT        # hop-by-hop options for ipv6
icmp    1       ICMP          # internet control message protocol
igmp    2       IGMP          # internet group management protocol
ggp     3       GGP           # gateway-gateway protocol
ipencap 4       IP-ENCAP      # IP encapsulated in IP (officially ``IP'')
st      5       ST            # ST datagram mode
tcp     6       TCP           # transmission control protocol
udp     17      UDP           # user datagram protocol
```

Common protocols include TCP, UDP and ICMP. There are actually a great number of protocols, but these three are the most prevalent. TCP is used to handle the underlying data streams for web, email, news and almost every other popular network application. That does not mean that UDP and ICMP are not important however. While TCP is good at providing reliable, error-free data UDP is used for applications where speed or simplicity is more important than reliability. Applications like DNS lookups and SNMP use the UDP protocol. ICMP is used for untilities such as **ping** and **traceroute**. If you have ever seen the `destination host unreachable` message, this is ICMP in action.

4. Network Configuration

Many modern networking environments use Dynamic Host Configuration Protocol (DHCP) [http://en.wikipedia.org/wiki/Dynamic_Host_Configuration_Protocol] to automatically set up the parameters needed to get computers connected to the network. However, a successful LPI candidate should understand how to manually configure Linux machines using the **ifconfig [http://www.die.net/doc/linux/man/man8/ifconfig.8.htmlhttp://www.die.net/doc/linux/man/man8/ifconfig.8.html]** and **route [http://www.die.net/doc/linux/man/man8/route.8.html]** commands. LPI exams generally ask straightforward questions, but it is not unusual to see a question asking about a common command-line option [http://www.lpi.org/en/faq2.html#2.14]. Be prepared by practicing manual network configuration with **ifconfig** and **route**.

5. Network Troubleshooting

A network of computers is a complex system and many times things do not work as well as intended. A network administrator needs to understand basic troubleshooting techniques like pinging and tools like netstat. (more to come)

6. PPP Configuration

under construction

7. Practice Questions

1. Given an IP address of 192.168.12.17 and a subnet mask of 255.255.255.0, what is the network portion of the address?

2. Given the CIDR address format of 172.16.10.1/24, what is the network mask in dotted-quad notation?

    A. 255.0.0.0

    B. 255.255.0.0

    C. 255.255.255.0

    D. 255.255.255.255

3. During a freak electrical storm both the primary and secondary DNS servers on your network were simultaneously struck by lightning and are no longer operational. In the absense of name servers which file can be used to do domain name to IP address lookups? (give the full path)

4. The loopback address, 127.0.0.1, is which class of IP address?

    A. Class A

    B. Class B

    C. Class C

    D. Class D

5. Your company's Internet connection is down and you have been called in to investigate. You would like to know if the fault lies with one of the routers on your LAN or if there is a problem with a router on your ISP's network. Which command would you use to quickly determine the location of problem?

6. An entry such as the following is found in which file? (give the complete path)

    ```
    smtp      25/tcp     # simple mail transport protocol
    ```

7. Several of your company's employees have asked for the ability to check their work email from home via the Internet. You have configured an IMAP daemon to accomodate them. On which port does IMAP communicate?

    A. 23

    B. 25

    C. 110

    D. 143

8. The DHCP server for your LAN has a failed power supply and it will take 24 hours for the new part to arrive. Which command can be used to manually configure IP addresses until the DHCP server can be repaired?

    A. netstat

    B. ipconfig

    C. ifconfig

    D. inetcfg

9. In an effort to maintain a well-documented network environment you want to do periodic checks of network activity on your server. Which command will let you view active network connections? (specify the command without any command-line options.)

8. Answers To Practice Questions

1. The correct answer is 192.168.12.0.

2. Answer C is correct, /24 in CIDR notation means a 24-bit mask. Each octet in an IP address has 8 bits and the decimal equivalent of eight bits all set to 1 is 255. So answer A is equivalent to /8 in CIDR notation, answer B is the same as /16 and answer D is /32.

3. The correct answer is `/etc/hosts`.

4. The correct answer is Class A, any address with a first octet that is less than 128 is considered to be a Class A address.

5. The correct answer is **traceroute**. While you might be tempted to use **ping**, this would require you to ping each router individually and does not *quickly* determine where the problem is.

6. The correct answer is `/etc/services`.

7. The IMAP protocol uses port 143 so D is the correct answer. Ports 23, 25 and 110 are used for telnet, smtp and pop3, respectively.

8. Answer C, **ifconfig** is correct. **netstat** is used to view network connections, not configure interfaces. **ipconfig** and **inetcfg** are both commands used by other operating systems.

9. The correct answer is **netstat**.

9. Additional References

under construction

Chapter 15. Networking Services

1. This Chapter Is Up For Adoption

If you feel you are an expert in the subject area for this LPI topic and would like to write this chapter please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide to communicate your intentions. Authors should be LPI certified or hold a similar certification in a related subject.

Please adopt only one chapter at a time and plan to finish the chapter within thirty days of adopting it.

More information about the study guide can be found at: http://www.happy-monkey.net/LPI/

1.1. Structuring The Chapter

The structure of chapters within the study guide needs to be consistent. The following structure may be used as a guide. Authors are also encouraged to view the chapter covering "Hardware & Architecture" as a template.

1. Author Name

    This is to ensure that authors get credit for their work. Please include your level of LPI certification after your name. If there are other appropriate certifications, please include these as well. For example, someone writing a chapter on networking might include the fact that they have a Cisco certification in addition to the LPI certification like this: "Joe Smith, LPIC-1, CCNA".

2. A Brief Look At The Objectives

    Each chapter should start with a high-level look at the LPI testing objectives. Within this "Brief Look" section there should be hyperlinks to sources of additional information. For example, one might say, "The networking section of the exam requires familiarity requires familiarity with IP addresses and subnet masks." The phrases 'IP addresses' and 'subnet masks' should be hyperlinked to additional information. Wikipedia is the prefered source of external information, but it is also helpful to link to man pages for specific commands.

3. Detailed Discussion

    Each one of the LPI exam objectives should receive a section for more detailed coverage of the objective and how it relates to Linux and the LPI exam. A good way to do this is by using hands-on exercises and examples from real-world Linux systems whenever possible. For example, when discussing a configuration file in /etc, give a sample of what a typical file might look like and encourage the reader to **cat** the file on their own system.

    Objectives may be grouped together when it is logical to do so.

4. Practice Questions

    At the end of each chapter there should be a section dedicated to practice questions that cover the exam objectives. *These should not be verbatim questions from the exam!* The idea is to let readers test their mastery of the knowledge in each chapter and get a feel for the exam format, not to help them cheat. Anyone intentionally submitting verbatim questions will be reported to the LPI and risks losing their certification.

1.2. DocBook XML Markup

The LPI Study Guide is written using DocBook XML markup. This allows the guide to be published in HTML, PDF and other formats using the same source document. Those familiar with DocBook are encouraged to submit chapters this way. Be sure to use markup for commands and filenames in addition to sections. Use the chapter on "Hardware

& Architecture" as a template.

For those not familiar with DocBook should submit chapters in plain text in a format similar to the example below.

```
A Brief Look At The Objectives

Successful completion of the Hardware & Architecture section
of the LPI exam requires familiarity with Personal Computer
(PC) expansion cards [http://url-to-wikipedia/expansion-cards]...

Next-Section-Title

More text...
```

1.3. Additional Questions

Please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide with any other questions not covered.

Chapter 16. Security

1. This Chapter Is Up For Adoption

If you feel you are an expert in the subject area for this LPI topic and would like to write this chapter please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide to communicate your intentions. Authors should be LPI certified or hold a similar certification in a related subject.

Please adopt only one chapter at a time and plan to finish the chapter within thirty days of adopting it.

More information about the study guide can be found at: http://www.happy-monkey.net/LPI/

1.1. Structuring The Chapter

The structure of chapters within the study guide needs to be consistent. The following structure may be used as a guide. Authors are also encouraged to view the chapter covering "Hardware & Architecture" as a template.

1. Author Name

    This is to ensure that authors get credit for their work. Please include your level of LPI certification after your name. If there are other appropriate certifications, please include these as well. For example, someone writing a chapter on networking might include the fact that they have a Cisco certification in addition to the LPI certification like this: "Joe Smith, LPIC-1, CCNA".

2. A Brief Look At The Objectives

    Each chapter should start with a high-level look at the LPI testing objectives. Within this "Brief Look" section there should be hyperlinks to sources of additional information. For example, one might say, "The networking section of the exam requires familiarity requires familiarity with IP addresses and subnet masks." The phrases 'IP addresses' and 'subnet masks' should be hyperlinked to additional information. Wikipedia is the prefered source of external information, but it is also helpful to link to man pages for specific commands.

3. Detailed Discussion

    Each one of the LPI exam objectives should receive a section for more detailed coverage of the objective and how it relates to Linux and the LPI exam. A good way to do this is by using hands-on exercises and examples from real-world Linux systems whenever possible. For example, when discussing a configuration file in /etc, give a sample of what a typical file might look like and encourage the reader to **cat** the file on their own system.

    Objectives may be grouped together when it is logical to do so.

4. Practice Questions

    At the end of each chapter there should be a section dedicated to practice questions that cover the exam objectives. *These should not be verbatim questions from the exam!* The idea is to let readers test their mastery of the knowledge in each chapter and get a feel for the exam format, not to help them cheat. Anyone intentionally submitting verbatim questions will be reported to the LPI and risks losing their certification.

1.2. DocBook XML Markup

The LPI Study Guide is written using DocBook XML markup. This allows the guide to be published in HTML, PDF and other formats using the same source document. Those familiar with DocBook are encouraged to submit chapters this way. Be sure to use markup for commands and filenames in addition to sections. Use the chapter on "Hardware

& Architecture" as a template.

For those not familiar with DocBook should submit chapters in plain text in a format similar to the example below.

```
A Brief Look At The Objectives

Successful completion of the Hardware & Architecture section
of the LPI exam requires familiarity with Personal Computer
(PC) expansion cards [http://url-to-wikipedia/expansion-cards]...

Next-Section-Title

More text...
```

1.3. Additional Questions

Please email the author [mailto:dhorton.no-spam(at)no-spam.member.fsf.org] of the LPI study guide with any other questions not covered.