

The Linux Plug-and-Play-HOWTO

Table of Contents

<u>The Linux Plug-and-Play-HOWTO</u>	1
<u>David S.Lawyer mailto:dave@lafn.org</u>	1
<u>1.Introduction</u>	1
<u>2.What PnP Should Do: Allocate "Bus-Resources"</u>	1
<u>3.The Plug-and-Play (PnP) Solution</u>	1
<u>4.Configuring a PnP BIOS</u>	2
<u>5.How to Deal with PnP Cards</u>	2
<u>6.Tell the Driver the Configuration</u>	2
<u>7.What Is My Current Configuration?</u>	2
<u>8.Appendix</u>	2
<u>1. Introduction</u>	3
<u>1.1 Copyright, Trademarks, Disclaimer, & Credits</u>	3
<u>Copyright</u>	3
<u>Disclaimer</u>	3
<u>1.2 Future Plans: You Can Help</u>	3
<u>1.3 New Versions of this HOWTO</u>	4
<u>2. What PnP Should Do: Allocate "Bus-Resources"</u>	4
<u>2.1 What is Plug-and-Play (PnP)?</u>	4
<u>2.2 How a Computer Finds Devices (and conversely)</u>	4
<u>2.3 I/O Addresses, etc.</u>	5
<u>2.4 IRQs --Overview</u>	6
<u>2.5 DMA Channels</u>	6
<u>2.6 Memory Ranges</u>	7
<u>2.7 "Resources" to both Device and Driver</u>	7
<u>2.8 The Problem</u>	7
<u>2.9 PnP Finds Devices Plugged Into Serial Ports</u>	8
<u>3. The Plug-and-Play (PnP) Solution</u>	8
<u>3.1 Introduction to PnP</u>	8
<u>3.2 How It Works (simplified)</u>	9
<u>3.3 Starting Up the PC</u>	10
<u>3.4 Buses</u>	10
<u>3.5 Linux Needs to Cope Better with PnP</u>	10
<u>4. Configuring a PnP BIOS</u>	11
<u>4.1 Do you have a PnP operating system?</u>	11
<u>Interoperability with Windows9x</u>	12
<u>4.2 How are bus-resources to be controlled?</u>	12
<u>4.3 Reset the configuration?</u>	12
<u>5. How to Deal with PnP Cards</u>	13
<u>5.1 Introduction to Dealing with PnP Cards</u>	13
<u>5.2 Disable PnP ?</u>	13
<u>5.3 BIOS Configures PnP</u>	14
<u>Intro to Using the BIOS to Configure PnP</u>	14
<u>The BIOS's ESCD Database</u>	14
<u>Using Windows to set the ESCD</u>	15
<u>Adding a New Device (under Linux or Windows)</u>	16
<u>5.4 Isapnp (part of isapnptools)</u>	16
<u>5.5 PCI Utilities</u>	17

Table of Contents

5.6 Patch the Kernel to Make Linux PnP.....	17
5.7 Windows Configures	17
5.8 Device Driver Configures.....	17
5.9 PnP Software/Documents.....	17
6. Tell the Driver the Configuration.....	18
6.1 Introduction.....	18
6.2 Serial Port Driver: setserial.....	18
6.3 Sound Card Drivers.....	19
OSS-Lite.....	19
OSS (Open Sound System) and ALSA.....	19
7. What Is My Current Configuration?.....	19
7.1 How Are My Device Drivers Configured?.....	20
7.2 How Are My Hardware Devices Configured?.....	20
8. Appendix.....	21
8.1 Addresses.....	21
ISA Bus Configuration Address (Read-Port etc.).....	21
Address ranges.....	21
Address space.....	22
Range Check (ISA Testing for IO Address Conflicts).....	22
Communicating Directly via Memory.....	22
8.2 Interrupts --Details.....	22
8.3 PCI Interrupts.....	23
8.4 Isolation.....	24

The Linux Plug-and-Play-HOWTO

David S.Lawyer <mailto:dave@lafn.org>

v0.09, 7 January 2000

Help with understanding and dealing with the complex Plug-and-Play issue. How to get your Linux system to support Plug-and-Play.

1.Introduction

- [1.1 Copyright, Trademarks, Disclaimer, & Credits](#)
- [1.2 Future Plans: You Can Help](#)
- [1.3 New Versions of this HOWTO](#)

2.What PnP Should Do: Allocate "Bus-Resources"

- [2.1 What is Plug-and-Play \(PnP\)?](#)
- [2.2 How a Computer Finds Devices \(and conversely\)](#)
- [2.3 I/O Addresses, etc.](#)
- [2.4 IRQs ---Overview](#)
- [2.5 DMA Channels](#)
- [2.6 Memory Ranges](#)
- [2.7 "Resources" to both Device and Driver](#)
- [2.8 The Problem](#)
- [2.9 PnP Finds Devices Plugged Into Serial Ports](#)

3.The Plug-and-Play (PnP) Solution

- [3.1 Introduction to PnP](#)
- [3.2 How It Works \(simplified\)](#)
- [3.3 Starting Up the PC](#)
- [3.4 Buses](#)
- [3.5 Linux Needs to Cope Better with PnP](#)

4. Configuring a PnP BIOS

- [4.1 Do you have a PnP operating system?](#)
- [4.2 How are bus-resources to be controlled?](#)
- [4.3 Reset the configuration?](#)

5. How to Deal with PnP Cards

- [5.1 Introduction to Dealing with PnP Cards](#)
- [5.2 Disable PnP ?](#)
- [5.3 BIOS Configures PnP](#)
- [5.4 Isapnp \(part of isapnptools\)](#)
- [5.5 PCI Utilities](#)
- [5.6 Patch the Kernel to Make Linux PnP](#)
- [5.7 Windows Configures](#)
- [5.8 Device Driver Configures](#)
- [5.9 PnP Software/Documents](#)

6. Tell the Driver the Configuration

- [6.1 Introduction](#)
- [6.2 Serial Port Driver: setserial](#)
- [6.3 Sound Card Drivers](#)

7. What Is My Current Configuration?

- [7.1 How Are My Device Drivers Configured?](#)
- [7.2 How Are My Hardware Devices Configured?](#)

8. Appendix

- [8.1 Addresses](#)
- [8.2 Interrupts --Details](#)
- [8.3 PCI Interrupts](#)
- [8.4 Isolation](#)

[Next](#) [Previous](#) [Contents](#) [Next](#) [Previous](#) [Contents](#)

1. Introduction

1.1 Copyright, Trademarks, Disclaimer, & Credits

Copyright

Copyright (c) 1998, 1999 by David S. Lawyer.

Please freely copy and distribute (sell or give away) this document in any format. Forward any corrections and comments to the document maintainer. You may create a derivative work and distribute it provided that you:

1. Send your derivative work (in the most suitable format such as sgm) to the LDP (Linux Documentation Project) or the like for posting on the Internet. If not the LDP, then let the LDP know where it is available. Except for a translation, send a copy to the previous maintainer's url as shown in the latest version.
2. License the derivative work in the spirit of this license or use GPL. Include a copyright notice and at least a pointer to the license used.
3. Give due credit to previous authors and major contributors.

If you're considering making a derived work other than a translation, it's requested that you discuss your plans with the current maintainer.

Disclaimer

Much of the info in this HOWTO was obtained from the Internet, implications in books that may be in error or obsolete, etc. While I haven't intentionally tried to mislead you, there are likely a number of errors in this document. Please let me know about them. Since this is free documentation, it should be obvious that I cannot be held legally responsible for any errors.

1.2 Future Plans; You Can Help

Please let me know of any errors in facts, opinions, logic, spelling, grammar, clarity, links, etc. But first, if the date is over a month old, check to see that you have the latest version. Please send me any info that you think belongs in this document.

I haven't studied in detail neither isapnptools nor David Howells' patches to the kernel (but I plan to). Nor do I fully understand how PnP is configured by the BIOS (it depends on which BIOS) nor how Windows9x updates the ESCD. Thus this HOWTO is still incomplete and may be inaccurate (let me know where I'm wrong). In this HOWTO I've sometimes used ?? to indicate that I don't really know the answer.

1.3 New Versions of this HOWTO

New versions of the Plug-and-Play-HOWTO should appear every month or so and will be available to browse and/or download at LDP mirror sites. For a list of mirror sites see: <http://metalab.unc.edu/LDP/mirrors.html>. Various formats are available. If you only want to quickly check the date of the latest version look at: <http://metalab.unc.edu/LDP/HOWTO/Plug-and-Play-HOWTO.html>. The version you are now reading is: v0.08, 28 November 1999.

[Next](#) [Previous](#) [Contents](#)[Next](#)[Previous](#)[Contents](#)

2. What PnP Should Do: Allocate "Bus-Resources"

2.1 What is Plug-and-Play (PnP)?

Oversimplified, Plug-and-Play automatically tells the software (device drivers) where to find various pieces of hardware (devices) such as modems, network cards, sound cards, etc. Plug-and-Play's task is to match up physical devices with the software (device drivers) that operates them and to establish channels of communication between each device and its driver. In order to achieve this, PnP allocates the following "bus-resources" to both drivers and hardware: I/O addresses, IRQ's, DMA channels (ISA bus only), and memory regions. If you don't understand what these 4 bus-resources are read the following subsections of this HOWTO. An article in Linux Gazette about 3 of these bus-resources is [Introduction to IRQs, DMAs and Base Addresses](#) Once these bus-resources have been assigned (and if the correct driver is installed), the names for such devices in the /dev directory are ready to use.

This PnP assignment of bus-resources is sometimes called "configuring" but it is only a low level type of configuring. Even with PnP fully utilized, much configuring of devices is done by other than PnP. For example, for modem configuration an "init string" is sent to the modem over the I/O address "channel". This "init string" has nothing to do with PnP although the "channel" used to send it to the modem was allocated by PnP. Setting the speed (and many other parameters) of a serial port is done by sending messages to the device driver from programs run by the user (often automatically boot-time). This configuring also has nothing to do with PnP. Thus when talking about PnP "configuring" means only a certain type of configuring. While other documentation (such a for MS Windows) simply calls bus-resources "resources", I have coined the term "bus-resources" so as to distinguish it from the multitude of other kinds of resources.

2.2 How a Computer Finds Devices (and conversely)

A computer consists of a CPU/processor to do the computing and memory to store programs and data. In addition, there are a number of devices such as various kinds of disk-drives, a video card, a keyboard, network cards, modem cards, sound cards, serial and parallel ports, etc. There is also a power supply to provide electric energy, various buses on a motherboard to connect the devices to the CPU, and a case to put all this into.

The Linux Plug-and-Play-HOWTO

In olden days most all devices had their own plug-in cards (printed circuit boards). Today, in addition to plug-in cards, many "devices" are small chips permanently mounted on the "motherboard". Cards which plug into the motherboard may contain more than one device. Memory chips are also sometimes considered to be devices but are not plug-and-play in the sense used in this HOWTO.

For the computer system to work right, each device must be under the control of its "device driver". This is software which is a part of the operating system (perhaps loaded as a module) and runs on the CPU. Device drivers are associated with "special files" in the /dev directory although they are not really files. They have names such as hda1 (first partition on hard drive a), ttyS0 (the first serial port), eth1 (the second ethernet card), etc. To make matters more complicated, the particular device driver selected, say for eth1, will depend on the type of ethernet card you have. Thus eth1 can't just be assigned to any ethernet driver. It must be assigned to a certain driver that will work for the type of ethernet card you have installed. To control a device, the CPU (under the control of the device driver) sends commands (and data) to and reads info from the various devices. In order to do this each device driver must know the address of the device it controls. Knowing such an address is equivalent to setting up a communication channel, even though the physical "channel" is actually the data bus inside the PC which is shared with almost everything else.

The communication channel is actually a little more complex than described above. An "address" is actually a range of addresses and there is a reverse part of the channel (known as interrupts) which allows devices to send an urgent "help" request to their device driver.

2.3 I/O Addresses, etc.

PC's have 3 address spaces: I/O, main memory, and configuration (only on the PCI bus). All of these 3 types of addresses share the same address bus inside the PC. But the voltage on certain dedicated wires on the PC's bus tells which "space" an address is in: I/O, main memory, or configuration. See [Addresses](#) for more details. Devices were originally located in I/O address space although today they may use space in main memory. An I/O address is sometimes just called "I/O", "IO", "i/o" or "io". The term "I/O port" also used. There are two main steps to allocate the I/O addresses (or other bus-resources such as interrupts):

1. Set the I/O address, etc. on the card (in one of its registers)
2. Let its device driver know what this I/O address, etc. is

The two step process above is something like the two part problem of finding someone's house number on a street. You must obtain (and write down) the house number and someone must install a number on the front of the house so that it may be found. In computers, the device driver must obtain the address and the device hardware must get the same address set in one of its registers. Both of these must be done, but some people make the mistake of doing only one of these and then wonder why the computer can't find the device. For example, they will use "setserial" to assign an address to a serial port without realizing that this only tells the driver an address. It doesn't set the address in the serial port hardware itself. If the serial port actually had a different address (or none at all) and you told setserial wrong, then you're in trouble.

Another obvious requirement is that before the device driver can use an address it must be first set on the card. Since device drivers often start up soon after you start the computer, they sometimes try to access a card (to see if it's there, etc.) before the address has been set in the card by a PnP configuration program. Then you see an error message that they can't find the card even though it's there (but doesn't yet have an address).

What was said in the last 2 paragraphs regarding I/O addresses applies with equal force to other

bus-resources: [IRQs --Overview](#), [DMA Channels](#), and [Memory Regions](#). What these are will be explained in the next 3 sections.

2.4 IRQs --Overview

After reading this you may read [Interrupts --Details](#) for some more details. The following is intentionally oversimplified: Besides the address, there is also an interrupt number to deal with (such as IRQ 5). It's called an IRQ (Interrupt ReQuest) number. We already mentioned above that the device driver must know the address of a card in order to be able to communicate with it. But what about communication in the opposite direction? Suppose the the device needs to tell its device driver something immediately? For example, the device may have just received a lot of bytes destined for main memory and the device needs to call its driver to fetch these bytes at once and transfer them from the device's nearly full buffer into main memory.

How should the device call for help? It can't use the main data bus since it's likely already in use. Instead it puts a voltage on an dedicated interrupt wire (part of the bus) which is often reserved for that device alone. This signal is called an interrupt. There are the equivalent of 16 such wires in a PC and each wire leads (indirectly) to a certain device driver. Each wire has a unique IRQ (Interrupt ReQuest) number. The device must put its interrupt on the correct wire and the device driver must listen for the interrupt on the correct wire. Which wire it's put on is determined by the IRQ number stored in the device. This same IRQ number must be known to the device driver so that the device driver knows which IRQ line to listen to.

Once the device driver gets the interrupt (a call for help) it must find out why the interrupt was issued and take appropriate action to service the interrupt. On the ISA bus each device needs its own unique IRQ number. For the PCI bus and other special cases the sharing of IRQs is allowed.

2.5 DMA Channels

DMA channels are only for the ISA bus. DMA stands for "Direct Memory Access". This is where a device is allowed to take over the main computer bus from the CPU and transfer bytes directly to main memory. Normally the CPU would make such a transfer in a two step process: 1. reading from the I/O memory space of the device and putting these bytes into the CPU itself 2. writing these bytes from the CPU to main memory. With DMA it's usually a one step process of sending the bytes directly from the device to memory. The device must have such capabilities built into its hardware and thus not all devices can do DMA. While DMA is going on the CPU can't do too much since the main bus is being used by The DMA transfer.

The PCI bus doesn't really have any DMA but instead it has something even better: bus mastering. It works something like DMA and is sometimes called DMA (for example, hard disk drives that call themselves "UltraDMA"). It allows devices to temporarily become bus masters and to transfer bytes almost like the bus master was the CPU. It doesn't use any channel numbers since the organization of the PCI bus is such that the PCI hardware knows which device is currently the bus master and which device is requesting to become a bus master. Thus there is no allocation of DMA channels for the PCI bus.

When a device on the ISA bus wants to do DMA it issues a DMA-request using dedicated DMA request wires much like an interrupt request. DMA actually could have been handled by using interrupts but this would introduce some delays so it's faster to do it by having a special type of interrupt known as a DMA-request. Like interrupts, DMA-request are numbered so as to identify which device is making the request. This number is called a DMA-channel. Since DMA transfers all use the main bus (and only one can

run at a time) they all actually use the same channel but the "DMA channel" number serves to identify who is using the "channel". Hardware registers exist on the motherboard which store the current status of each "channel". Thus in order to issue a DMA-request, the device must know its DMA-channel number which must be stored in a register on the physical device.

2.6 Memory Ranges

Some devices are assigned address space in main memory. It's often "shared memory" or "memory-mapped I/O". Sometimes it's ROM memory on the device. When discussing bus-resources it's often just called "memory". Such a device might also use I/O address space.

When you plug in such a card, you are in effect also plugging in a memory module for main memory. This memory can either be ROM (Read Only Memory) or shared memory. Shared memory is shared between the device and the CPU (running the device driver). This memory can serve as a means of direct data "transfer" between the device and main memory. It's not really a transfer since the device puts data into its own memory on its card which also happens to be in main memory. Both the card and the device driver need to know where it is. The memory address are likely to be very high so that they do not conflict with the lower addresses of the memory chips in your computer.

ROM is different. It is likely a program (perhaps a device driver) which will be used with the device. Hopefully, it may work with Linux and not just Windows ?? It may need to be shadowed which means that it is copied to your main memory chips in order to run faster. Once it's shadowed it's no longer "read only".

2.7 "Resources" to both Device and Driver

Thus device drivers must be "attached" in some way to the hardware they control. This is done by supplying bus-resources (I/O, Memory, IRQ's, DMA's) to both the physical device and the device driver software. For example, a serial port uses only 2 (out of 4 possible) resources: an IRQ and an I/O address. Both of these values must be supplied to the device driver and the physical device. The driver (and its device) is also given a name in the /dev directory (such as ttyS1). The address and IRQ number is stored by the physical device in registers on the card (or in a chip on the motherboard). For the case of jumpers, this info is always stored in the device hardware (on the card, etc.). But for the case of PnP, the register data is usually lost when the PC is powered down (turned off) so that the resource data must be supplied to each device anew each time the PC is powered on.

2.8 The Problem

The architecture of the PC provides only a limited number of IRQ's, DMA channels, I/O address, and memory regions. If there were only several devices and they all had standardized bus-resource (such as unique I/O addresses and IRQ numbers) there would be no problem of attaching device drivers to devices. Each device would have a fixed resources which would not conflict with any other device on your computer. No two devices would have the same addresses, there would be no IRQ conflicts, etc. Each driver would be programmed with the unique addresses, IRQ, etc. hard-coded into the program. Life would be simple.

But it's not. Not only are there so many different devices today that conflicts are frequent, but one sometimes needs to have more than one of the same type of device. For example, one may want to have a few different disk-drives, a few serial ports, etc. For these reasons devices need to have some flexibility so that they can be set to whatever address, IRQ, etc. is needed to avoid conflicts. But some IRQ's and addresses are pretty standard such as the ones for the clock and keyboard. These don't need such flexibility.

Besides the problem of conflicting allocation of bus-resources, there is a problem of making a mistake in telling the device driver what the bus-resources are. For example, suppose that you enter IRQ 4 in a configuration file when the device is actually set at IRQ 5. This is another type of bus-resource allocation error.

The allocation of bus-resources, if done correctly, establishes channels of communication between physical hardware and their device drivers. For example, if a certain I/O address range (resource) is allocated to both a device driver and a piece of hardware, then this has established a one-way communication channel between them. The driver may send commands and info to the device. It's actually a little more than one-way since the driver may get information from the device by reading it's registers. But the device can't initiate any communication this way. To initiate the device needs an IRQ in order to create a two-way communication channel where both the driver and the device can initiate communication.

2.9 PnP Finds Devices Plugged Into Serial Ports

External devices that connect to the serial port via a cable (such as external modems) can also be called Plug-and-Play. Since only the serial port itself needs bus-resources (an IRQ and I/O address) there are no bus-resources to allocate to such plug-in devices. Thus PnP is not really needed for them. Even so, there is a PnP specification for such external serial devices.

A PnP operating system will find such an external device and read its model number, etc. Then it may be able to find a device driver for it so that you don't have to tell an application program that you have a certain device on say /dev/ttyS1. Since you should be able to manually inform your application program (via a configuration file, etc.) what serial port the device is on (and possibly what model number it is) you should not really need this "serial port" feature of PnP.

[NextPreviousContentsNextPreviousContents](#)

3. The Plug-and-Play (PnP) Solution

3.1 Introduction to PnP

The term Plug-and-Play (PnP) has various meanings. In the broad sense it is just auto-configuration where one just plugs in a device and it configures itself. In the sense used in this HOWTO, the configuration is only that of configuring PnP bus-resources and letting the device drivers know about it. In a more narrow sense it just setting bus-resources in the hardware devices. It may also mean the PnP specifications which (among

other things) specify how PnP resource data is to be read and written to devices (often cards) on the ISA bus. The standard PCI (and not PnP) specifications do the same for the PCI bus.

PnP matches up devices with their device drivers and specifies their communication channels. On the ISA bus before Plug-and-Play the bus-resources were set in hardware devices by jumpers. Software drivers were assigned bus-resources by configuration files (or the like) or by probing the for the device at addresses where it's expected to reside. The PCI bus was PnP-like from the beginning so it was trivial to implement PnP for this bus. Since the PCI bus specifications don't use the term PnP it's not clear whether or not the PCI bus should be called PnP (but it supports in hardware what today is called PnP).

3.2 How It Works (simplified)

Here's an oversimplified view of how PnP works. The PnP configuration program (perhaps a program in the BIOS) finds all PnP devices and asks each what bus-resources it needs. Then it checks what bus-resources (IRQs, etc.) it has to give away. Of course if it has reserved bus-resources used by non-PnP (legacy) devices (if it knows about them) it doesn't give these away. Then it uses some criteria (not specified by PnP specifications) to give out the bus-resources so that there are no conflicts and so that all devices get what they need (if possible). It then tells each physical device what bus-resources are assigned to it and the devices set themselves up to use only the assigned bus-resources. Then the device drivers somehow find out what bus-resources their devices use and are thus able to communicate effectively with the devices they control.

For example, suppose a card needs one interrupt (IRQ number) and 1 MB of shared memory. The PnP program reads this request from the card. It then assigns the card IRQ5 and 1 MB of memory addresses space, starting at address 0xe9000000. It's not always this simple as the card may specify that it can only use certain IRQ numbers (ISA only) or that the 1 MB of memory must lie within a certain range of addresses. The details are different for the PCI and ISA buses with more complexity on the ISA bus.

There are some shortcuts that PnP software may use. One is to keep track of how it assigned bus-resources at the last configuration (when the computer was last used) and reuse this. Windows9x and PnP BIOSs do this but standard Linux doesn't. Windows9x stores this info in its "Registry" on the hard disk and a PnP BIOS stores it in non-volatile memory in your PC (known as ESCD; see [The BIOS's ESCD Database](#)).

Under Linux it's each device for itself and there is no centralized non-volatile registry of resource assignments. Some device drivers store the last configuration they used and use it next time the computer is powered on. They implicitly assume that the rest of the hardware will not need to use it's bus-resources.

If the device hardware remembered their previous configuration, then there wouldn't be any hardware to configure at the next boot-time, but they seem to forget their configuration when the power is turned off. Some devices contain a default configuration (but not necessarily the last one used). Thus a PnP configuration program needs to be run each time the PC is powered on. Also, if a new device has been added, then it too needs to be configured. Allocating bus-resources to this new device might involve taking some bus-resources away from an existing device and assigning the existing device alternative bus-resources that it can use instead.

3.3 Starting Up the PC

When the PC is first turned on the BIOS chip runs its program to get the computer started (the first step is to check out the hardware). If the operating system is stored on the hard-drive (as it normally is) then the BIOS must know about the hard-drive. If the hard-drive is PnP then the BIOS may use PnP methods to find it. Also, in order to permit the user to manually configure the BIOS's CMOS and respond to error messages when the computer starts up, a screen (video card) and keyboard are also required. Thus the BIOS must PnP-configure these devices on its own.

Once the BIOS has identified the hard-drive, the video card, and the keyboard it is ready to start booting (loading the operating system into memory from the hard-disk). If you've told the BIOS that you have a PnP operating system (PnP OS), it should start booting the PC as above and let the operating system finish the PnP configuring. Otherwise, a PnP-BIOS will (prior to booting) likely try to do the rest of the PnP configuring of devices (but not their drivers).

3.4 Buses

ISA is the old bus of the old IBM PC's while PCI is a newer and faster bus from Intel. The PCI bus was designed for what is today called PnP. It makes it easy (as compared to the ISA bus) to find out how PnP bus-resources have been assigned to hardware devices. To see what has happened look at the /proc/pci "file" (/proc/bus/pnp/devices for kernel 2.2+), the boot-up messages on your display (use shift-PageUp to back up), or use PCI Utilities (for kernel 2.2+).

For the ISA bus there is a real problem with implementing PnP since no one had PnP in mind when the ISA bus was designed and there are almost no I/O addresses available for PnP to use for sending configuration info to physical device. As a result, the way PnP was shoehorned onto the ISA bus is very complicated. A whole book has been written about it. See [PnP Book](#). Among other things, it requires that each PnP device be assigned a temporary "handle" by the PnP program so that one may address it for PnP configuring. Assigning these "handles" is called "isolation". See [Isolation](#) for the complex details.

Eventually, the ISA bus should become extinct. When it does, PnP will be easier since it will be easy to find out how the BIOS has configured the hardware. There will still be the need to match up device drivers with devices and also a need to configure devices that are added when the PC is up and running. These needs would be satisfied if Linux was a PnP operating system.

3.5 Linux Needs to Cope Better with PnP

PnP (for the ISA bus) was invented by Compaq, Intel, and Phoenix. Microsoft has been a leading promoter of it. Linux would have been better off if PnP had never been "invented". Eventually the ISA bus will have become extinct and the PnP-like PCI bus will prevail so that we will have in effect gotten an easy-to-implement PnP. But like it or not, most all new ISA hardware today is PnP and Linux has no choice but to deal effectively with PnP. But standard Linux (as of early 1999) makes dealing with PnP complicated (especially on the ISA bus) while the purpose of PnP was to make it simple.

In a sense, Linux is already somewhat PnP for the PCI bus. When the PC starts up you may note from the messages on the screen that some Linux device drivers often find their hardware devices (and the

bus-resources the BIOS has assigned them). But there are situations that a PnP operating system could handle better: 1. A shortage of bus-resources 2. More than one driver for a physical device 3. An activated driver which can't find it's physical device 4. Hot installation of a device (docking, etc.)

Linux users should not need to delve into the details of PnP to configure ISA PnP devices as they now need to. One solution would be a standardized version of the Linux kernel that supports Plug-and-Play on the ISA, PCI, and other buses. A patch to the kernel has been written although most drivers don't support it. It's not part of standard Linux. See [Patch Kernel](#).

[Next](#)[Previous](#)[Contents](#)[Next](#)[Previous](#)[Contents](#)

4. Configuring a PnP BIOS

When the computer is first turned on, the BIOS runs before the operating system is loaded. Newer BIOSs are PnP and will configure some or all of the PnP devices. For most PnP BIOSs there is no way to disable PnP so you have to live with it. Here are some of the choices which may exist in your BIOS's CMOS menu:

- [Do you have a PnP operating system?](#)
- [How are bus-resources to be controlled?](#)
- [Reset the configuration?](#)

4.1 Do you have a PnP operating system?

If you say yes, then the PnP BIOS will PnP-configure the hard-drive, video card, and keyboard to make the system bootable. But it will leave it up to the operating system to finish the configuration job. It may do an [Isolation](#) on the ISA bus leaving the devices disabled but ready to be configured by the operating system. For Linux you should probably tell it that you don't have a PnP operating system. If you don't do this, the BIOS might leave the ISA devices it hasn't configured in a disabled state ?? Also PCI devices might not get configured ??

If you tell the BIOS you don't have a PnP OS, then the BIOS will do the configuring itself. Unless you have added new PnP devices, it should use the configuration which it has stored in its non-volatile memory (ESCD). See [The BIOS's ESCD Database](#) If the last session on your computer was with Linux, then there should be no change in configuration. See [BIOS Configures PnP](#). But if the last session was with Windows9x (which is PnP) then Windows could have modified the ESCD. It supposedly does this only if you "force" a configuration or install a legacy device. See [Using Windows to set ESCD](#). If you are using the isapnp or PCI Utilities program(s) to do configuring, they will run after the BIOS runs and change things the way you told them to.

Interoperability with Windows9x

If you are running both Linux and Windows on the same PC, how do you answer the BIOS's question: Do you have a PnP OS? Normally (and truthfully) you would say no for standard Linux and yes for Windows9x. But it's a lot of bother to have to set up the BIOS's CMOS menu manually each time you want to switch OSs. One solution is set the CMOS for no PnP OS, including when you start Windows. One might expect that Windows would be able to handle this situation where it is presented hardware that has been fully configured by the BIOS. In addition, one might expect that even if Windows didn't realize that the hardware was already configured, it would redo this configuration and then work OK. But it doesn't seem to work this way. It seems that Windows may just tell its device drivers what has been stored in the Windows' Registry. But the actual hardware configuration (done by the BIOS) is what was stored in the ESCD and may not be the same as the Registry => trouble.

One way to try to get the Registry and the ESCD the same is to install (or reinstall) Windows when the BIOS is set for "not a PnP OS". This should present Windows with hardware configured by the BIOS. If this configuration is without conflicts, Windows will hopefully leave it alone and save it in it's Registry. Then the ESCD and the registry are in sync. If this works for you (and this is the latest version of this HOWTO), let me know as I only have one report of this working out OK.

Another method is to "remove" devices that are causing problems in Windows by clicking on remove in the Device Manager. Then reboot with "Not a PnP OS" (set it in the CMOS as you start to boot). Windows will then reinstall the devices, hopefully using the bus-resource settings configured by the BIOS. Be warned that Windows will likely ask you to insert the Windows installation CD since it sometimes can't find the driver files (and the like) even though they are still there. As a test I "removed" the NIC card which had a Novell compatible drivers. Upon rebooting, Windows reinstalled it with Microsoft networking instead of Novell. This meant that the Novell Client needed to be reinstalled. Let me know about your problems with this method (only if this is the latest version of this HOWTO).

4.2 How are bus-resources to be controlled?

This may involve just deciding how to allocate IRQ and DMA bus-resources. If set to "auto", the BIOS will do the allocation. If set to manual, you manually reserve some IRQ's for use on "legacy" (non-pnp) cards. The BIOS may or may not otherwise know about your legacy cards. The BIOS will only know about your legacy cards if you ran ICU (or the like) under Windows to tell the BIOS about them. If the BIOS knows about them, then try using "auto". If it doesn't know about them then manually reserve the IRQ's needed for the legacy ISA cards and let the rest be for the BIOS PnP to allocate.

4.3 Reset the configuration?

This will erase the BIOS's ESCD data-base of how your PnP devices should be configured as well as the list of how legacy (non-PnP) devices are configured. Never do this unless you are convinced that this data-base is wrong and needs to be remade. It was stated somewhere that you should do this only if you can't get your computer to boot. If the BIOS loses the data on legacy devices, then you'll need to run ICA again under DOS/Windows to reestablish this data.

5. How to Deal with PnP Cards

5.1 Introduction to Dealing with PnP Cards

Today most all new internal boards (cards) are Plug-and-Play (PnP). Although some software exists in Linux to handle PnP, it is not always easy to use. There are 6 different methods listed below to cope with PnP (but some may not be feasible in your situation). Which one(s) you should use depends on your goals. What may be most expedient to do now may not be the easiest and best in the long run. A seemingly simple way is to do nothing and just let a PnP-BIOS configure it but then you may need to do some exploring to find out what the BIOS has done. A comparison of these methods needs to be written by someone who has tried them all. You may need to use more than one method to do the job.

- [Disable PnP](#) by jumpers or DOS/Windows software (but many cards can't do this)
- [BIOS Configures PnP](#) (For the PCI bus you only need a PCI BIOS, otherwise you need a PnP BIOS)
- [Isapnp](#) is a program you can always use to configure PnP devices on the ISA bus only
- [PCI Utilities](#) is for configuring the PCI bus
- [Windows Configures](#) and then you boot Linux from within Windows/DOS. Use as a last resort
- [Patch Kernel](#) to transform Linux into a PnP operating system
- [Device Driver Configures](#) but few do

Any of the above will set the bus-resources in the hardware. But only the last two should tell device driver what it's done. Only the last one definitely tells the driver (since it is the driver). How the driver gets informed depends on the driver and you may need to do something to inform it. See [Tell the Driver the Configuration](#)

5.2 Disable PnP ?

Many devices are PnP only with no option for disabling PnP. But for some, you may be able to disable PnP by jumpers or by running a Windows program that comes with the device (jumperless configuration). This will avoid the often complicated task of configuring PnP. Don't forget to tell the BIOS that these bus-resources are reserved. There are also some reasons why you might not want to disable PnP:

1. If you have MS Windows on the same machine, then you may want to allow PnP to configure devices differently under Windows from what it does under Linux.
2. The range of selection for IRQ numbers (or port addresses) etc. may be quite limited unless you use PnP.
3. You might have a Linux device driver that uses PnP methods to search for the device it controls.
4. If you need to change the configuration in the future, it may be easier to do this if it's PnP (no setting of jumpers or running a Dos/Windows program).

5. You may have (or will have) other PnP devices that need configuring so that you'll need to provide for (or learn about) PnP anyway.

Once configured as non-PnP devices, they can't be configured by PnP software or the BIOS (until you move jumpers and/or use the Dos/Windows configuration software again).

5.3 BIOS Configures PnP

Intro to Using the BIOS to Configure PnP

If you have a PnP BIOS, it can configure the hardware. This means that your BIOS reads the resource requirements of all devices and configures them (allocates bus-resources to them). It is a substitute for a PnP OS except that the BIOS doesn't match up the with their devices nor tell the drivers how it has done the configuring. It should normally use the configuration it has stored in its non-volatile memory (ESCD). If it finds a new device or if there's a conflict, the BIOS should make the necessary changes to the configuration and will not use exactly what was in the ESCD.

Your BIOS must support such configuring but there have been cases where it doesn't do it correctly or completely. An advantage of using the BIOS is that it's simple since in most cases there is nothing to set up (except to tell the BIOS's CMOS menu it's not a PnP OS). While some device drivers may be able to automatically detect what the BIOS has done, in some cases you'll need to determine it (not always easy). See [What Is My Current Configuration?](#) Another possible advantage is that the BIOS does its work before Linux starts so that all the bus-resources are ready to be used (and found) by the device drivers that start up later.

According to MS it's only optional (not required) that a PnP BIOS be able to PnP-configure the devices (without help from MS Windows). But it seems that most of the ones made after 1996 ?? or so can do it. We should send them thank-you notes if they do it right. They configure both the PCI and ISA buses, but it has been claimed that some older BIOSs can only do the PCI. To try to find out more about your BIOS, look on the Web. Please don't ask me as I don't have data on this. The details of the BIOS that you would like to know about may be hard to find (or not available). Some BIOSs may have minimal PnP capabilities and try to turn over the difficult parts of the configuration task to Window utilities. If this happens you'll either have to find another method (such as isapnptools) or try to set up the ESCD database if the BIOS has one. See the next section.

The BIOS's ESCD Database

The BIOS's maintains a non-volatile database containing a PnP-configuration that it will try to use. It's called the ESCD (Extended System Configuration Data). Again, the provision of ESCD is optional but most PnP-BIOSs have it. The ESCD not only stores the resource-configuration of PnP devices but also stores configuration information of non-PnP devices (and marks them as such) so as to avoid conflicts. The ESCD data is usually saved on a chip and remains intact when the power is off, but sometimes it's kept on a hard-drive??

The ESCD is intended to hold the last used configuration, but if you use a program such as Linux's isapnp or pci utilities (which doesn't update the ESCD) then the ESCD will not know about this and will not save this

The Linux Plug-and-Play-HOWTO

configuration in the ESCD. A good PnP OS might update the ESCD so you can use it later on for a non-PnP OS (like standard Linux). MS Windows does this only in special cases. See [Using Windows to set ESCD](#).

To use what's set in ESCD be sure you've set "Not a PnP OS" or the like in the BIOS's CMOS. Then each time the BIOS starts up (before the Linux OS is loaded) it should configure things this way. If the BIOS detects a new PnP card which is not in the ESCD, then it must then allocate bus-resources to the card and update the ESCD. It may even have to change the bus-resources assigned to existing PnP cards and modify ESCD accordingly.

If each device saved its last configuration in its hardware, hardware configuring wouldn't be needed each time you start your PC. But it doesn't work this way. So all the ESCD data needs to be kept correct if you use the BIOS for PnP. There are some BIOSs that don't have an ESCD but do have some non-volatile memory to store info on which bus-resources have been reserved for use by non-PnP cards. Many BIOSs have both.

Using Windows to set the ESCD

If the BIOS doesn't set up the ESCD the way you want it (or the way it should be) then it would be nice to have a Linux utility to set the ESCD. As of early 1999 there isn't any. Thus one may resort to attempting to use Windows (if you have it on the same PC) to do this.

There are three ways to use Windows to try to set/modify the ESCD. One way is to use the ICU utility designed for DOS or Windows 3.x. It should also work OK for Windows 9x/2k ?? Another way is to set up devices manually ("forced") under Windows 9x/2k so that Windows will put this info into the ESCD when Windows is shut down normally. The third way is only for legacy devices that are not plug-and-play. If Windows knows about them and what bus-resources they use, then Windows should put this info into the ESCD.

If PnP devices are configured automatically by Windows without the user "forcing" it to change settings, then such setting will probably not make it into the ESCD. Of course Windows may well decide on its own to configure the same as what is set in the ESCD so they could wind up being the same by coincidence.

Windows 9x are PnP operating systems and automatically PnP-configure devices. They maintain their own PnP-database deep down in the Registry (stored in binary Windows files). There is also a lot of other configuration stuff in the Registry besides PnP-bus-resources. There is both a current PnP resource configuration in memory and another (perhaps about the same) stored on the hard disk. To look at (the one in memory?) this indirectly in Windows98 or to force changes you use the Device Manager.

In Windows98 There are 2 ways to get to the Device Manager: 1. My Computer --> Control Panel --> System Properties --> Device Manager. 2. (right-click) My Computer --> Properties --> Device Manager. Then in Device Manager you select a device (sometimes a multi-step process if there are a few devices of the same class). Then click on Properties and then on Resources. To attempt to change the resource configuration manually, uncheck "Use automatic settings" and then click on "Change Settings". Now try to change the setting, but it may not let you change it. If it does let you, you have "forced" a change. A message should inform you that it's being forced. If you want to keep the existing setting shown by Windows but make it "forced" then you will have to force a change to something else and then force it back to its original setting.

To see what has been "forced" under Windows98 look at the "forced hardware" list: Start --> Programs --> Accessories --> System Tools --> System Information --> Hardware Resources --> Forced Hardware. When you "force" a change of bus-resources in Windows, it should put your change into the ESCD

(provided you exit Windows normally). >From the "System Information" window you may also inspect how IRQs and IO ports have been allocated under Windows.

Even if Windows shows no conflict of bus-resources, there may be a conflict under Linux. That's because Windows may assign bus-resources differently than the ESCD does. In the the rare case where all devices under Windows are either legacy devices or have been "forced", then Windows and the ESCD configurations should be identical.

Adding a New Device (under Linux or Windows)

If you add a new PnP device and have the BIOS set to "not a PnP OS", then the BIOS should automatically configure it and store the configuration in ESCD. If it's a non-PnP legacy device (or one made that way by jumpers, etc.) then there are a few options to handle it.

You may be able to tell the BIOS directly (via the CMOS setup menus) that certain bus-resources it uses (such as IRQs) are reserved and are not to be allocated by PnP. This does not put this info into the ESCD. But there may be a BIOS menu selection as to whether or not to have these CMOS choices override what may be in the ESCD in case of conflict. Another method is to run ICU under DOS/Windows. Still another is to install it manually under Windows 9x/2k and then make sure its configuration is "forced" (see the previous section). If it's "forced" Windows should update the ESCD when you shut down the PC.

5.4 Isapnp (part of isapnptools)

Unfortunately, much of the documentation for isapnp is still difficult to understand unless you know the basics of PnP. This HOWTO should help you understand it as well the FAQ that comes with it. `isapnp` is only for PnP devices on the ISA bus (non-PCI). Running the Linux program "isapnp" at boot-time will configure such devices to the resource values specified in `/etc/isapnp.conf`. Its possible to create this configuration file automatically but you then must edit it manually to chose between various options. With isapnp, a device driver which is part of the kernel may run too early before isapnp has set the address, etc. in the hardware. This results in the device driver not being able to find the device. The driver trys the right address but the address hasn't been set yet in the hardware.

If your Linux distribution automatically installed isapnptools, isapnp may already be running at startup. In this case, all you need to do is to edit `/etc/isapnp.conf` per "`man isapnp.conf`". Note that this is like manually configuring PnP since you make the decisions as to how to configure as you edit the configuration file. You can use the program "pnpdump" to help create the configuration file. It almost creates a configuration file for you but you must skillfully edit it a little before using it. It contains some comments to help you edit it. If you use "isapnp" for configuring and have a PnP BIOS, you should probably tell the BIOS (when you set it up) that you don't have a PnP OS since you may want the BIOS to configure the PCI devices. While the BIOS may also configure the ISA devices, isapnp will redo it.

The terminology used in the `/etc/isapnp.conf` file may seem odd at first. For example for an IO address of 0x3e8 you might see "(IO 0 (BASE 0x3e8))" instead. The "IO 0" means this is the first (0th) IO address-range that this device uses. Another way to express all this would be: `IO[0] = 0x3e8`" but isapnp doesn't do it this way. "IO 1" would mean that this is the second IO address range used by this device, etc. "INT 0" has a similar meaning but for IRQs (interrupts). A single card may contain several physical devices and the above explanation was for just one of these devices.

5.5 PCI Utilities

The new package PCI Utilities (= pciutils, incorrectly called "pcitools"), should let you manually PnP-configure the PCI bus. "lspci" lists bus-resources while "setpci" sets resource allocations in the hardware devices.

5.6 Patch the Kernel to Make Linux PnP

David Howells has created a patch to do this called "Linux Kernel Configuration/Resource Manager" (sometimes called Hardware Configuration Manager). In late 1999 the patch was not available at his website. This may mean there is no patch available for recent versions of the kernel.

For previous patches the resulting kernel was claimed to be stable but bugs have been reported. The patch included documentation such as serial.txt to show how to deal with the serial port. It provided "files" in the /proc tree so that you can see what is going on and can echo commands into one of these files for custom configuration. One problem is that most device drivers don't know about it so that you still had to use the traditional configuration files, etc. for configuration. The webpage for it is

<http://www.astarte.free-online.co.uk>

5.7 Windows Configures

If you have Windows9x (or 2k) on the same PC, then just start Windows and let it configure PnP. Then start Linux from Windows (or DOS). It been reported that Windows erased the IRQs from PCI devices registers. Then Linux complained that it found a zero IRQ. Thus you may not be able to use this method.

5.8 Device Driver Configures

A few device drivers will use PnP methods to set the bus-resources in the hardware but only for the device that they control. Since the driver has done the configuring, it obviously knows the configuration and there is no need for you to tell it this info.

The problem with this is twofold. It's difficult to incorporate all of this into the driver, and the driver may grab bus-resources that are needed by other devices. It does make it easy for the user but a PnP Linux kernel might be better. See [Linux Needs to Cope Better with PnP](#)

5.9 PnP Software/Documents

- [Isapnptools homepage](#)
- [Patch to make the Linux kernel PnP](#)
- [PnP driver project](#)
- [PnP Specs. from Microsoft](#)

- Book: PCI System Architecture, 3rd ed. by Tom Shanley +, MindShare 1995. Covers PnP-like features on the PCI bus.
- Book: Plug and Play System Architecture, by Tom Shanley, Mind Share 1995. Details of PnP on the ISA bus. Only a terse overview of PnP on the PCI bus.

[NextPreviousContentsNextPreviousContents](#)

6. Tell the Driver the Configuration

6.1 Introduction

Just how this is done depends upon the driver. Some drivers have more than one way to find out how their physical device is configured. At one extreme is the case where you must hard-code the bus-resources into the kernel and recompile. At the other extreme, the driver does everything automatically and you have nothing to do. It may even set the bus-resources in the hardware using PnP methods.

In the middle are cases where you run a program to give the resource info to the driver or put the info in a file. In some cases the driver may probe for the device at addresses where it suspects the device resides. It may then try to test various IRQs to see which one works. It may or may not automatically do this. In other cases the driver may use PnP methods to find the device and how the bus-resources have been set, but will not actually set them. It may also look in some of the files in the /proc directory.

One may need to give the bus-resources as a parameter to the kernel to a loadable module. See /usr/lib/modules_help/descr.gz for a list of possible parameters. The module to load is listed in /etc/modules along with its parameters. In some other case the bus-resources may be given as parameters to the kernel. These are put into the lilo.conf file as append="...". Then the lilo program must be run to save this in the kernel boot code.

While there is great non-uniformity about how drivers find out about bus-resources, the end goal is the same. There are so many different hardware devices and drivers for them that you may need to look at documentation for your driver to find out how it finds out about bus-resources and what you need to do to insure that it gets the info it needs. Some brief info on a few drivers is presented in the following section.

6.2 Serial Port Driver: setserial

For the standard serial port driver (not for multiport cards) you use setserial to configure the driver. It is often run from a start-up file. In newer versions there is a /etc/serial.conf file that you "edit" by simply using the setserial command in the normal way and what you set using setserial is saved in the serial.conf configuration file. The serial.conf file should be consulted when the setserial command runs from a start-up file. Your distribution may or may not set this up for you.

There are two different ways to use setserial depending on the options you give it. One way is used to

manually tell the driver the configuration. The other way is to probe at a given address and report if a serial port exists there. It can also probe this address and try to detect what IRQ is used for this port. The driver runs something like `setserial` at start-up but it doesn't probe for IRQs, it just assigns the "standard" IRQ which may be wrong. It does probe for the existence of a port. See `Serial-HOWTO` for more details.

6.3 Sound Card Drivers

OSS-Lite

You must give the IO, IRQ, and DMA as parameters to a module or compile them into the kernel. But some PCI cards will get automatically detected (likely by using the `lspci` command or the like). RedHat supplies a program "sndconfig" which detects ISA PnP cards and automatically sets up the modules for loading with the detected bus-resources.

OSS (Open Sound System) and ALSA

These will detect the card by PnP methods and then select the appropriate driver and load it. It will also set the bus-resources on an ISA-PnP card. You may need to manually intervene to avoid conflicts. For the ALSA driver, support for ISA-PnP is optional and you may use `isapnp` tools if you want to.

[NextPreviousContentsNextPreviousContents](#)

7. What Is My Current Configuration?

Here "configuration" means the assignment of PnP bus-resources (addresses, IRQs, and DMAs). There are two parts to this question for each device. Each part should have the same answer.

1. What is the configuration of the device driver software? I.e.: What does the driver think the hardware configuration is?
2. What configuration (if any) is set in the device hardware?

Of course the configuration of the device hardware and its driver should be the same (and it normally is). But if things are not working right, there may be a difference. This means the the driver has incorrect information about the actually configuration of the hardware. This spells trouble. If the software you use doesn't adequately tell you what's wrong (or automatically configure it correctly) then you need to investigate how your hardware devices and their drivers are configured. While Linux device drivers should "tell all" in some cases it's not easy to determine what has been set in the hardware.

Another problem is that when you view configuration messages on the screen, it's sometimes not clear whether the reported configuration is that of the device driver, the device hardware, or both. If the device driver is assigned a configuration and then checks the hardware out to see if it's configured the same, then the

configuration reported by the driver should be that of both the hardware and the driver.

But some drivers don't do this may accept a configuration that doesn't check out. For example, "setserial" will accept a configuration that doesn't check out (even if you've told it to probe for bus-resources). Thus "setserial" may only be telling you the configuration of the driver and not the hardware.

Some info on configuration may be obtained by reading the messages from the BIOS and Linux that appear on the screen when you first start the computer. These messages often flash by too fast to read but once they stop type Shift-PageUp a few times to scroll back to them. To scroll forward thru them type Shift-PageDown. Typing "dmesg" at any time to the shell prompt will show only the Linux kernel messages and miss some of the most important ones (including ones from the BIOS). The messages from Linux may sometimes only show what the device driver thinks the configuration is, perhaps as told it via an incorrect configuration file.

The BIOS messages will show the actual hardware configuration at that time, but a PnP OS, isapnp, or pci utilities, may change it later. The BIOS messages are displayed first before the ones from Linux. As an alternative to eventually using Shift-PageUp to read them, try freezing them by hitting the "Pause" key. Press any key to resume. But once the messages from Linux start to appear, it's too late to use "Pause" since it will not freeze the messages from Linux.

7.1 How Are My Device Drivers Configured?

There may be a programs you can run from the command line (such as "setserial" for serial ports) to determine this. The /proc directory tree is useful. /proc/ioports shows the I/O addresses that the drivers use (or try if it's wrong). They might not be set this way in hardware.

/proc/interrupts shows only interrupts currently in use and many that have been allocated to drivers don't show at all since they're not currently being used. For example, even though your floppy drive has a floppy disk in it and is ready to use, the interrupt for it will not show unless its in use. Again, just because an interrupt shows up here doesn't mean that it exists in the hardware. A clue that it doesn't exist in hardware will be if it shows that 0 interrupts have been issued by this interrupt. Even if it shows some interrupts have been issued, it may mean that this interrupt doesn't exist on that device does exist on some other device which is not in use, but which somehow has issued an interrupt or two. As of kernel 2.2 the /proc tree has changed.

7.2 How Are My Hardware Devices Configured?

It's easy to find out what bus-resources have been assigned to devices on the PCI bus with the "lspci" command. For for kernels <: 2.2: /proc/pci will provide some info but it's not too easy to decipher. For the ISA bus you may try running `pnpdump --dumpregs` but it's not a sure thing. The results may be difficult to decipher. Don't confuse the read-port address which `pnpdump` "trys" (and finds something there) with the I/O address of the found device. They are not the same.

Messages from the BIOS at boot-time tell you how the hardware configuration was then. If you rely on the BIOS for configuring, then it should still be the same. Messages from Linux may be from drivers that have checked to see that the hardware is there (and possibly checked the IRQ and DMA). Of course, if the device works fine, then it's likely configured the same as the driver.

8. Appendix

8.1 Addresses

There are three types of addresses: main memory addresses, I/O addresses and configuration addresses. On the PCI bus, configuration addresses constitute a separate address space just like I/O addresses do. Except for the complicated case of ISA configuration addresses, whether or not an address on the bus is a memory address, I/O address, or configuration address depends only on the voltage on other wires (traces) of the bus.

ISA Bus Configuration Address (Read-Port etc.)

For the ISA bus, there is technically no configuration address space, but there is a special way for the CPU to access PnP configuration registers on the PnP cards. For this purpose 3 @ I/O addresses are allocated. This is not 3 addresses for each card but 3 addresses shared by all cards.

These 3 addresses are named read-port, write-port, and address-port. Each port is just one byte in size. Each PnP card has many configuration registers so that just 3 addresses are not even sufficient for these registers on a single card. To communicate with a certain card, a specially-assigned card number (handle) is sent to all cards at the write-port address. After that only that the only card still listening is the card with this handle. Then the address of the configuration register (of that card) is sent to the address-port (of all cards --but only one is listening). Next communication takes place with one configuration register on that card by either doing a read on the read-port or a write on the write-port.

The write-port is always at A79 and the address-port is always at 279 (hex). But the read-port is not fixed but is set by the configuration software at some address that will supposedly not conflict with any other ISA card. If there is a conflict, it will change the address. All PnP cards get "programmed" with this address. Thus if you use say isapnp to set or check configuration data it must determine this read-port address.

Address ranges

The term "address" is sometimes used in this document to mean a contiguous range of addresses. Since addresses are given in bytes, a single address only contains one byte but I/O (and main memory) addresses need more than this. So a range of say 8 bytes is often used for I/O address while the range for main memory addresses allocated to a device is much larger. For a serial port (an I/O device) it's sufficient to give the starting I/O address of the device (such as 3F8) since it's well known that the range of addresses for serial port is only 8 bytes. The starting address is known as the "base address".

Address space

For ISA, to access both I/O and (main) memory address "spaces" the same address bus is used (the wires used for the address are shared). How does the device know whether or not an address which appears on the address bus is a memory address or I/O address? Well, there are 4 dedicated wires on the bus that convey this information and more. If a certain one of these 4 wires is asserted, it says that the CPU wants to read from an I/O address, and the main memory ignores the address on the bus. The other 3 wires serve similar purposes. In summary: Read and write wires exist for both main memory and I/O addresses (4 wires in all).

For the PCI bus it's the same basic idea also using 4 wires but it's done a little differently. Instead of only one or the four wires being asserted, a binary number is put on the wires (16 different possibilities). Thus more info may be conveyed. Four of these 16 numbers serve the I/O and memory spaces as in the above paragraph. In addition there is also configuration address space which uses up two more numbers. Ten extra numbers are left over for other purposes.

Range Check (ISA Testing for IO Address Conflicts)

On the ISA bus, there's a method built into each PnP card for checking that there are no other cards that use the same address. If two or more cards use the same IO address, neither card is likely to work right (if at all). Good PnP software should assign bus-resources so as to avoid this conflict, but even in this case a legacy card might be lurking somewhere with the same address.

The test works by a card putting a test number in its own IO registers. Then the PnP software reads it and verifies that it reads the same test number. If not, something is wrong (such as another card with the same address). It repeats the same test with another test number. Since it actually checks the range of IO addresses assigned to the card, it's called a "range check". It could be better called an address-conflict test. If there is an address conflict you get an error message and need to resolve it yourself.

Communicating Directly via Memory

Traditionally, most I/O devices used only I/O memory to communicate with the CPU. For example, the serial port does this. The device driver, running on the CPU would read and write data to/from the I/O address space and main memory. A faster way would be for the device itself to put the data directly into main memory. One way to do this is by using [DMA Channels](#) or bus mastering. Another way is to allocate some space in main memory to the device. This way the device reads and writes directly to main memory without having to bother with DMA or bus mastering. Such a device may also use IO addresses.

8.2 Interrupts --Details

Interrupts convey a lot of information but only indirectly. The interrupt signal (a voltage on a wire) just tells a chip called the interrupt controller that a certain device needs attention. The interrupt controller then signals the CPU. The CPU finds the driver for this device and runs a part of it known as an "interrupt service routine" (or "interrupt handler"). This "routine" tries to find out what has happened and then deals with the problem such as transferring bytes from (or to) the device. This program (routine) can easily find out what has happened since the device has registers at addresses known to the driver software (provided the IRQ

number and the I/O address of the device has been set correctly). These registers contain status information about the device. The software reads the contents of these registers and by inspecting the contents, finds out what happened, and takes appropriate action..

Thus each device driver needs to know what interrupt number (IRQ) to listen to. On the PCI bus (and for the serial ports on the ISA bus starting with Kernel 2.2) it's possible for two (or more) devices to share the same IRQ number. When such an interrupt is issued, the CPU runs all interrupt service routines for all devices using that interrupt. The first thing the first service routine does is to check to see if an interrupt actually happened for its device. If there was no interrupt (false alarm) it likely will exit and the next service routine starts, etc.

8.3 PCI Interrupts

PCI interrupts are different but since they are normally mapped to IRQ's they behave in about the same way. A major difference is that PCI interrupts may be shared. For example IRQ5 may be shared between two PCI devices. This sharing ability is automatic: you don't need special hardware or software. There have been some reports of situations where such sharing didn't work, but it's likely due to a defect in the device driver software. All device drivers for PCI devices are supposed to provide for interrupt sharing. Note that you can't share the same interrupt between the PCI and ISA bus. However, illegal sharing will work provided the devices which are in conflict are not in use at the same time. "In use" here means that a program is running which "opened" the device in it's C programming code.

You may need to know some of the details of the PCI interrupt system in order to set up the BIOS's CMOS or to set jumpers on old PCI cards. Each PCI card has 4 possible interrupts: INTA#, INTB#, INTC#, INTD#. Thus for a 7-slot system there could be $7 \times 4 = 28$ different interrupt lines. But the specs permit a fewer number of interrupt lines. This is not too restrictive since interrupts may be shared. Many PCI buses seem to be made with only 4 interrupt lines. Call these lines (wires or traces) W, X, Y, Z. Suppose we designate the B interrupt from slot 3 as interrupt 3B. Then wire W could be used to share interrupts 1A, 2B, 3C, 4D, 5A, 6B, 7C. This is done by physically connecting wire W to wires 1A, 2B, etc. Likewise wire X could be connected to wires 1B, 2C, 3D, 4A, 5B, 6C, 7D. Then on startup, the BIOS maps the X, W, Y, Z to IRQ's. After that it writes the IRQ that each device is mapped to into a hardware register in each device. Then and anything interrogating the device can find out what IRQ it uses.

The above mentioned wires X, W, Y, Z are labeled per PCI specs as INTA#, INTB#, INTC# and INTD#. This official PCI notation is confusing since now INTA# has 2 possible meanings depending on whether we are talking about a slot or the PCI bus. For example, if 3C is mapped to X then we say that INTC# of slot 3 is cabled to INTA# (X) of the PCI bus. Confusing notation.

There's another requirement also. A PCI slot must use the lower interrupt letters first. Thus if a slot only uses one interrupt, it must be INTA#. If it uses 2 interrupts they must be INTA# and INTB#, etc. A card in a slot may have up to 8 devices on it but there are only 4 PCI interrupts for it. This is OK since interrupts may be shared so that each of the 8 devices (if they exist) can have an interrupt. The PCI interrupt letter of a device is often fixed and hardwired into the device.

The BIOS assigns IRQs (interrupts) so as to avoid conflicts with the IRQs it knows about on the ISA bus. Sometimes in the CMOS BIOS menu one may assign IRQs to PCI cards (but it's not simple as explained above). There's a situation where Windows zeroed out all the IRQ numbers in the PCI cards after the IRQ mappings had been set. Then someone running Windows booted Linux from Windows with the result that Linux only found only incorrect IRQs of zero.

You might reason that since the PCI is using IRQ's (ISA bus) it might be slow, etc. Not really. The ISA Interrupt Controller Chip(s) has a direct interrupt wire going to the CPU so it can get immediate attention. While signals on the ISA address and data buses need to go thru the PCI bus to get to the CPU, the IRQ interrupt signals go there almost directly.

8.4 Isolation

This is only for the ISA bus. Isolation is a complex method of assigning a temporary handle (id number or Card Select Number = CSN) to each PnP device on the ISA bus. Since there are more efficient (but more complex) ways to do this, some might claim that it's a simple method. Only one write address is used for PnP writes to all PnP devices so that writing to this address goes to all PnP device that are listening. This write address is used to send (assign) a unique handle to each PnP device. To assign this handle requires that only one device be listening when the handle is sent (written) to this common address. All PnP devices have a unique serial number which they use for the process of isolation. Doing isolation is something like a game. It's done using the equivalent of just one common bus wire connecting all PnP devices and the isolation program.

For the first round of the "game" all PnP devices listen on this wire and send out simultaneously a sequence of bits to the wire. The allowed bits are either a 1 (positive voltage) or an "open 0" of no voltage (open circuit or tri-state). Each PnP device just starts to sequentially send out its serial number, bit-by-bit, starting with the high-order bit, on this wire. If any device sends a 1, a 1 will be heard on the wire by all other devices. If all devices send an "open 0" nothing will be heard on the wire. The object is to eliminate (by the end of this first round) all but highest serial number device. "Eliminate" means to cease to listen anymore to the write address that all devices still in the game are still listening to. This is also called "dropping out". (Note that all serial numbers are of the same length.)

First consider only the high order bit of the serial number which is put on the wire first by all devices which have no handle yet. If any PnP device sends out a 0 (open 0) but hears a 1, this means that some other PnP device has a higher serial number, so it temporarily drops out of this round and doesn't listen anymore until the round is finished (when a handle is assigned to the winner: the highest serial number). Now the devices still in the game all have the same leading digit (a 1) so we may strip off this digit and consider only the resulting "stripped serial number" for future participation in this round. Then go to the start of this paragraph and repeat until the entire serial number has been examined for each device (see below for the all-0 case).

Thus it's clear that the highest serial number will not be eliminated from the game. But what happens if the leading digits (of the possibly stripped serial numbers) are all 0? In this case an "open 0" is sent on the line and all participants stay in the game. If they all have a leading 0 then this is a tie and the 0's are stripped off just like the 1's were in the above paragraph. The game then continues as the next digit (of the serial number) is sent out.

At the end of the round (after the low-order bit of the serial number has been sent out by whatever participants remain) only one PnP device with the highest serial number remains. It then gets assigned a handle and drops out of the game permanently. Then all the dropouts from the last round (that don't have a handle yet) reenter the game and a new round begins with one less participant. Eventually, all PnP devices are assigned handles. It's easy to prove that this algorithm works.

Once all handles are assigned, they are used to address each PnP device and send it a configuration as well as to read configuration info from the PnP device. Note that these handles are only used for PnP configuration and are not used for normal communication with the PnP device. When the computer starts up, all of the

The Linux Plug-and-Play-HOWTO

handles are lost so that a PnP BIOS usually does the isolation process again each time you start your PC.

END OF Plug-and-Play-HOWTO

Next [PreviousContents](#)