
The Unix Hardware Buyer HOWTO

Eric Raymond <esr@thyrsus.com>

	Revision History	
Revision 4.2	2010-04-11	esr
	DVD region-locking firmware is no longer an issue,	
Revision 4.1	2009-07-01	esr
	DTX failed. Finally deprecate SCSI. 32-bit is dead. Avoiding the printer-consumables trap. Invasion of the netbooks.	
Revision 4.0	2007-11-02	esr
	Major revisions by Jonathan Marsden on SATA, bus standards, DVDs and other topics, followed by a cleanup pass from me.	
Revision 3.3	2007-18-13	esr
	Updated for 2007 conditions. CRTs are dead. BTX is dead. CD-ROMs are completely generic now. USB modems are recommended.	
Revision 3.2	2004-10-28	esr
	Fix and remove bad links.	
Revision 3.1	2004-08-03	esr
	Sound cards don't matter any more.	
Revision 3.0	2004-02-21	esr
	Power-protection stuff moved to UPS HOWTO. DIMM memory is gone. Tape drives don't make sense any more. Lots of the theory from my "Ultimate Linux Box" articles now lives here.	
Revision 2.4	2003-02-22	esr
	URL fixes.	
Revision 2.3	2002-08-06	esr
	Buying at the low end isn't a lose anymore. I recommend Athlons. Nuked the section on video standards, EDID takes care of all that now. Also removed the section on older memory types. And keyboards, as the "ergonomic" ones all vanished along with the 1990s carpal-tunnel scare.	
Revision 2.2	2002-08-05	esr
	New section on DVD drives.	
Revision 2.1	2002-07-08	esr
	Corrected Kingston URL. Various small updates for the last year. This HOWTO is much more stable than it used to be.	
Revision 2.0	2001-08-09	esr
	Major update. Revisions based on Ultimate Linux Box experience. Caches are on-chip now. DDS4 tape drives are here. 486 machines, CD caddies, and most non-DDS backup technologies are gone.	
Revision 1.1	2001-06-13	esr
	Mid-2001 update.	
Revision 1.0	2001-02-06	esr
	Initial revision; but see the history in the introduction.	

Abstract

This is your one-stop resource for information about how to buy and configure generic PC hardware for cheap, powerful Unix systems.

Table of Contents

Introduction	2
Purpose of this document	2
New versions of this document	3

Feedback and corrections	3
Related resources	4
Overview of the Market	4
Buying the Basics	4
Things to Not Care About	4
How To Pick Your Processor	6
One Disk or Two?	7
Getting Down to Cases	7
Power Supplies and Fans	8
Motherboards	9
Monitor and Video	9
DVD Drives	10
Sound Cards and Speakers	11
Modems	11
Printers	13
Power Protection	14
Radio Frequency Interference	14
What To Optimize	15
First, add more memory	15
Bus and Disk speeds	15
Optimizing your disk subsystem	16
Tuning Your I/O Subsystem	16
But What If I'm Economizing?	19
Noise Control and Heat Dissipation	20
Special Considerations When Buying Laptops and Netbooks	21
How to Buy	22
When to Buy	22
Where to Buy	23
Computer Fairs	23
Mail Order	23
Computer Superstores	24
Other Buying Tips	24
Questions You Should Always Ask Your Vendor	24
Minimum Warranty Provisions	24
Documentation	25
A System Quality Checklist	25
Things to Check when Buying	25
Tricks and Traps in Warranties	25
Special Questions to Ask Web/Mail-Order Vendors Before Buying	26
Payment Method	26
Which Clone Vendors to Talk To	26
After You Take Delivery	27
Software to go with your hardware	27
Other Resources on Building Linux PCs	27

Introduction

Purpose of this document

The purpose of this document is to give you the background information you need to be a savvy buyer of Intel hardware for running Unix. It is aimed especially at hackers and others with the technical skills and confidence to go to the Internet/mail-order channel, but contains plenty of useful advice for people buying store-front retail.

This document is maintained and periodically updated as a service to the net by Eric S. Raymond, who began it for the very best self-interested reason that he was in the market and didn't believe in plonking down several grand without doing his homework first (no, I don't get paid for this, though I have had a bunch of free software and hardware dumped on me as a result of it!). Corrections, updates, and all pertinent information are welcomed at esr@snark.thyrsus.com [<mailto:esr@snark.thyrsus.com>]. The editorial «we' reflects the generous contributions of many savvy Internetters.

If you email me questions that address gaps in the FAQ material, you will probably get a reply that says “Sorry, everything I know about this topic is in the HOWTO”. If you find out the *answer* to such a question, please share it with me for the HOWTO, so everyone can benefit.

If you end up buying something based on information from this HOWTO, please do yourself and the net a favor; make a point of telling the vendor “The HOWTO sent me” or some equivalent. If we can show vendors that this HOWTO influences a lot of purchasing decisions, we get leverage to change some things that need changing.

Note that in December 1996 I published an introductory article on building and tuning Linux systems summarizing much of the material in this HOWTO. It's available here [<http://www.catb.org/~esr/writings/lj-howtobuild.html>]. In 2001 I published an article on building the Ultimate Linux Box [<http://www.catb.org/~esr/writings/ultimate-linux-box/>].

This Buyer's Guide actually dates back to 1992, when it was known as the “PC-Clone Unix Hardware Buyer's Guide”; this was before Linux took over my world :-). Before that, portions of it were part of a Unix Buyer's Guide that I maintained back in the 1980s on USENET.

It may be a matter of historical interest that the page count of this guide peaked in mid-2001 and has been declining since. Video, sound, and other functions are migrating onto motherboards. Several bus types have disappeared, as have all the old-school backup technologies that couldn't scale up to match disk capacities, Spec sheets are getting simpler. Accordingly, there are parts that used to have whole sections to themselves that I barely even write about anymore — mice, floppy disks, CD-ROM drives, and keyboards, for example, are utterly generic now,

Another consequence of the technology stabilizing is also that I'm updating this guide less often than I used to. Years can now go by without the PC market changing in any fundamental way.

In retrospect, the success of the ATX standard for motherboards in 1998-1999 was probably the turning point. The PC industry has become sufficiently commoditized that your choices are now getting simpler rather than more complicated. This is a Good Thing.

New versions of this document

New versions of the Unix Hardware Buyer HOWTO will be periodically be uploaded to various Linux WWW and FTP sites, including the LDP home page.

You can view the latest version of this on the World Wide Web via the URL [HOWTO/Unix-Hardware-Buyer-HOWTO/](http://www.tldp.org/HOWTO/Unix-Hardware-Buyer-HOWTO/) [<http://www.tldp.org/HOWTO/Unix-Hardware-Buyer-HOWTO/>].

Feedback and corrections

If you have questions or comments about this document, please feel free to mail Eric S. Raymond, at esr@thyrsus.com [<mailto:esr@thyrsus.com>]. I welcome any suggestions or criticisms. If you find a mistake with this document, please let me know so I can correct it in the next version. Thanks.

Related resources

You may also want to look at the read the Hardware-HOWTO [<http://www.tldp.org/HOWTO/Hardware-HOWTO.html>]. It lists hardware known to be compatible with Linux, and hardware known to be incompatible. I've also done a series of articles on The Ultimate Linux Box [<http://www.catb.org/~esr/writings/ultimate-linux-box/>].

Overview of the Market

The central fact about PC hardware is that de-facto hardware standards have created a commodity market with low entry barriers, lots of competitive pressure, and volume high enough to amortize a *lot* of development on the cheap.

The result is that this hardware gives you lots of bang-per-buck, and it's getting both cheaper and better all the time. Furthermore, margins are thin enough that vendors have to be lean, hungry, and *very* responsive to the market to survive.

One good general piece of advice is that you should avoid the highest-end new-technology systems (those not yet shipping in volume). The problem with the high end is that it usually carries a hefty “prestige” price premium, and may be a bit less reliable on average because the technology hasn't been through a lot of test/improve cycles.

There used to be a real issue with low-end PCs as well, because there used to be a lot of dodgy crap PC components out there going into boxes made by vendors trying to save a few cents. That's not really a problem anymore. Market pressure has been very effective at raising reliability standards for even low-end components as the market has matured. It's actually hard to go wrong even buying at the bottom end of the market these days.

I put together the first version of this guide around 1992; Unix-capable systems are now ten to twenty times cheaper than they were then. At today's prices, building your own system from parts no longer makes much sense at all —so this HOWTO is now more oriented towards helping you configure a whole system from a single vendor.

Buying the Basics

In this section, we cover things to look out for that are more or less independent of price-performance tradeoffs, part of your minimum system for running Unix.

Issues like your choice of disk, processor, and I/O bus (where there is a significant tradeoff between price and capability) are covered in the section on What To Optimize.

Things to Not Care About

An effect of PC commoditization is that there aren't lots of things you used to have to worry about that don't matter any more, because the market has completely flattened out. We list these here to get them out of the way.

Bus Wars

The system bus is what ties all the parts of your machine together. This is an area in which progress has simplified your choices a lot. There used to be no fewer than *four* competing bus standards out there (ISA,

EISA, VESA/VLB, PCI, and PCMCIA). Now there are effectively just *two* —PCI-X on servers, and PCIe for desktop/tower machines. Even PCI is now legacy technology, and the PCMCIA bus that seemed so important a few years back has been reduced to near-irrelevance by Ethernet, USB, and WiFi hardware built onto motherboards. The newcomer is PCIe, which is (in late 2007) a ‘video-card-mostly’ bus, though it seems to be gaining in popularity for other uses too on mainstream desktop motherboards, whereas PCI-X is only found on higher end ‘server’ motherboards.

Memory

Judging the memory-controller and cache design used to be one of the trickiest parts of evaluating a motherboard, but that stuff is all baked into the processor itself now. This removed a large source of latency and design variations. It also killed off the plethora of different RAM types that used to be out there.

Today's advice is very simple. Make sure the memory is rated for your machine's bus speed, then buy as much as you can afford to stuff in your machine.

Note

DDR3 RAM is beginning to appear. Right now its extra expense over DDR2 is not worth paying, for all but extremely specialized needs. It is almost always *far* more useful to have 4GB of reasonably fast RAM, than 2GB of very fast RAM, in your machine.

For more technical stuff on memory architectures, see The Ultimate Memory Guide [<http://www.kingston.com/tools/umg/default.asp>] maintained by Kingston Technologies.

Keyboards and Mice

Keyboards are mostly generic nowadays. One useful piece of advice is to not buy any desktop machine with “Internet” buttons on it; this is a sure sign of a PC that's an overpriced glitzy toy. Nowadays keyboards with a USB connector are the norm, rather than the older dedicated connectors; modern open-source Unixes handle these just fine.

Mice and trackballs used to be simple; then, thanks to Microsoft, they got complicated. Now they're simple again. Again, USB mice have replaced the older PS/2-style dedicated connector. XFree86 autodetects your mouse when it starts up, so configuration is not a big deal any more.

Some PC vendors, being Windows-oriented, still bundle two-button mice. Thus, you may have to buy your own three-button (or two button and a scroll wheel) mouse. Ignore the adspeak about dpi and pick a mouse or trackball that feels good to your hand.

Your humble editor really, really likes the Logitech TrackMarble, an optical trackball that eliminates the chronic roller-fouling problems of the older TrackMan. They're well-supported by X, so any Linux or BSD will accept them.

Floppy Drives

There's no longer much to be said about floppy drives. They're cheap, they're generic, and the rise of CD-ROM and DVD-ROM drives as a cheap distribution medium has made them much less important than formerly. You only ever see the 3.5-inch ‘hard-shell’ floppies with 1.44MB capacity anymore.

Bootable CD-ROMs killed off the last use of floppies, which was OS installation. So go ahead and settle for cheap Mitsumi and Teac floppy drives. There are no ‘premium’ floppy drives anymore. Nobody bothers.

It's possible your system won't even include one. No loss.

CD-ROM Drives

Standard CD-ROMs hold about 650 megabytes of read-only data in a format called ISO-9660 (formerly “High Sierra”). All current Unixes support these devices. Unix and Linux software is now distributed on ISO-9660 CD-ROM, a cheaper and better method than the QIC tapes we used to use.

CD-ROM speed used to be a big deal; vendors advertised 2X, 4X, all the way up to 52X. Vendors don't bother any more; the drives are all about equivalently fast now.

There are one or two minor features to watch for. Most CD-ROMS will include a headphone jack so you can play audio CDs on them. Better-quality ones will also include two RCA jacks for use with speakers. Another feature to look for is a drive door or seal that protects the drive head from dust.

Increasingly, DVD-ROM drives (and burners) are replacing CD-ROM drives as the default optical drive in PC systems. They have significantly larger capacity, and will read (and burn) CD media too. The cost difference now is so small that it is usually preferable to buy a DVD burner instead of a CD-ROM drive.

Backup devices

It's good to be able to make backups that you can separate from your system and store off-site in case of disaster. Until about 2001, tape drives still seemed like a good idea for personal systems, but I found I seldom used mine. Today, tape drives with high enough capacity to image today's huge hard disks are too expensive to make sense any more.

For the money you'd spend on a high-capacity tape drive (over \$1000) it makes more sense to buy a laptop and a pile of CD-R or DVD-R or DVD+R media. Sit the laptop on your house Ethernet when you're not traveling, and back up the main machine to it every day, or oftener. Between the efficiency of rsync and the speed of 100-megabit Ethernet, this will be a lot faster than making a tape. Every once in a while, burn a set of backup CD-ROMs or DVDROMs.

But CD-ROMs aren't reusable; the cost piles up over time. An interesting alternative is a small external USB hard drive, especially if you can salvage an old laptop drive and put it in a USB enclosure. These enclosures are available for about \$30; Google for "USB HD Enclosure". This is faster than a tape, cheaper and lighter than a full laptop. For faster transfer speeds, an enclosure that accepts eSATA connections as well as USB helps a lot (assuming your PC or notebook has an eSATA connector).

How To Pick Your Processor

Right now (early 2010), the chips to consider for running Unix are the 64-bit AMD Opteron or its Intel equivalents, especially the Core 2 Duo. We're long past the point at which 32-bit chips are interesting for new desktop systems, presuming you could even find one. AMD and Intel built up a buffer before switching their fabs fully to 64-bit chips in 2006, and the 32-bit chips you can still find are coming out of warehouses rather than off production lines.

Brands don't matter much, so don't feel you need to pay Intel's premiums if you see an attractive Cyrix, AMD or other chip-clone system offered. In the last few years I've been a big fan of the AMD line. They used to be faster, cheaper, and better-designed than Intel processors; today Intel has clawed back the speed advantage, but AMD chips still deliver more performance than you're likely to be able to use and do it with lower power dissipation (thus, less noise and heat).

On the other hand, Intel-chip motherboards now have the advantage that the on-board graphics chip will give you 3D acceleration with fully open-source drivers. This will avoid the problems you would otherwise face trying to select a supported graphics card from ATI or Nvidia.

Many CPUs now are multi-core — that is, they have multiple CPUs on a single chip. This is very useful for doing something compute intensive (re-encoding video, compressing large archives, etc.) in the back-

ground and still having a responsive system for other work at the same time. At current prices, a dual-core CPU makes good sense for most desktop systems. If you are building a server or have specialized computing needs you expect to be very CPU-intensive quad-core is worth considering, but on a desktop system all the two extra cores will usually do is emit heat. Only at the very low end (sub US\$50 CPUs) do single-core CPUs still make sense on desktop machines.

Mainstream desktop CPUs now use one of two sockets: LGA 775 (Intel) and AM2 (AMD). Buying a system that uses one of these stands more chance of allowing a useful CPU upgrade to extend its useful life than systems using other less common sockets.

Current CPUs are *much* faster than those of just a few years ago. As a result, unless your needs are highly specialized, spending more than about US\$200 on a desktop CPU is hard to justify. For most users, putting extra budget into more RAM or a faster disk subsystem will most likely result in greater benefit.

One Disk or Two?

I usually build with two disks — one “system” disk and one “home” disk. There are two good reasons to do this that have nothing to do with the extra capacity. One of them is the performance advantage of being able to interleave commands to different physical spindles that we'll explain a bit later in the section on disks. The other is that I am quite a bit less likely to lose two disks at once than I am to trash a single one.

Let's suppose you have a fatal disk crash. If you have only one disk, goodbye Charlie. If you have two, maybe the crashed one was your system disk, in which case you can buy another and mess around with a new Linux installation knowing your personal files are safe. Or maybe it was your home disk; in that case, you can still run and do recovery stuff and basic Net communications until you can buy another home disk and restore it from backups (you *did* keep backups, right?).

Given today's high capacity drives, another way to use two disks well is to set them up as a RAID1 (mirrored) array. This can be done in software or with a hardware RAID controller. This way if either of the two drives fail, the system will continue to function, no data is lost, and upon replacing the failed drive, the array can be rebuilt from the remaining working drive. Hard drives are consumable media, they *do* fail, so this approach (as well as good backups) is well worth considering.

Buy SATA. The older IDE and EIDE buses are now obsolete, and SCSI no longer has enough of a cost advantage to justify the premium. In fact, SCSI has effectively merged into SATA; SATA is SCSI commands being shipped over a single-wire data line.

Getting Down to Cases

I used to say that cases are just bent metal, and that it doesn't much matter who makes those. Unfortunately, this isn't true any more. Processors run so hot these days that fans and airflow are a serious concern. They need to be well designed for proper airflow throughout.

Look for the following quality features:

- Aluminum rather than steel. It's lighter and conducts heat better.
- Unobstructed air intake with at least one fan each (in addition to the power supply and processor fans)
- No sharp metal edges. You don't want to shred your hands when you're tinkering with things.
- There shouldn't be any hot spots (poor air flow).
- Sturdy card clips. Some poorly-designed cases allow cards to wiggle out of their slots under normal vibration.

- Effective and easy to use mechanisms for attaching hard drives, CD-ROM, CD-R/W, DVDs, etc.

If you're fussy about RFI (Radio-Frequency Interference), it's worth finding out whether the plastic parts of the case have conductive coating on the inside; that will cut down emissions significantly, but a few cheap cases omit it.

Should you buy a desktop or tower case? Our advice is go with tower unless you're building a no-expansions personal system and expect to be using the floppies a lot. Many vendors charge nothing extra for a tower case, and the cost difference will be trivial even if they do. What you get for that is less desktop clutter, more and bigger bays for expansion, and often (perhaps most importantly) a beefed-up power-supply and fan. Putting the box and its fan under a table is good for maybe 5db off the effective noise level, too. Airflow is also an issue; if the peripheral bays are less cramped, you get better cooling. Be prepared to buy extension cables for your keyboard and monitor, though; vendors almost never include enough flex.

The airflow thing is a good argument for a full- or mid-tower rather than the 'baby tower' cases some vendors offer. However, smaller towers are getting more attractive as boards and devices shrink and more functions migrate onto the motherboard. A state of the art system, with all 3" disks, 300W power supply, half-size motherboard, on-board SATA and 4GB of RAM sockets, and half-sized expansion cards, will fit into a baby or midsized tower with ample room for expansion; and the whole thing will fit under a desk and make less noise than a classic tower.

For users with really heavy expandability requirements, rackmount PC cases do exist (ask prospective vendors). Typically a rackmount case will have pretty much the same functionality as an ordinary PC case. But, you can then buy drive racks (complete with power supply), etc. to expand into. Also, you can buy passive backplanes with up to 20 or so slots. You can either put a CPU card in one of the slots, or connect it to an ordinary motherboard through one of the slots.

Since USB has taken over most forms of detachable peripheral, a good feature to look for in a case is USB ports mounted at the top forward edge where it's easy to plug in digital cameras and the like.

Power Supplies and Fans

A lot of people treat power supplies as a commodity, so many interchangeable silver bricks. We know better — cheap power supplies go bad, and when they go bad they have a nasty habit of taking out the delicate electronics they're feeding. Also, the power supply tends to be the noisiest component in your system.

Give preference to supplies with a Underwriter's Laboratories rating. There's some controversy over optimum wattage level. On the one hand, you want enough wattage for expansion. On the other, big supplies are noisier, and if you draw too little current for the rating the delivered voltage can become unstable. And the expected wattage load from peripherals is dropping steadily. On the other hand, processors and their cooling fans eat a lot more power than they used to.

The choice is generally between 200W and 300W. After some years of deprecating 300W-and-up supplies as overkill, I'm now persuaded it's time to go back to them; a modern processor can consume 50-75W by itself, and for the newer dual-processor board the power supply needs to be rated 450W or up.

Processors on modern motherboards run hot enough that all vendors have gone to embedded temperature sensors and variable-speed thermostat-controlled fans, out sheer self-defense (this used to be a high-end only feature).

To cut noise, look for 120mm fans rather than the old-style 80mm muffin fans. These can move the same amount of air per minute rotating at a lower tip speed, which means less vortex formation and less noise. These are now becoming standard even on cheap white-box hardware.

In garden-variety tower cases there often isn't enough airflow to cool all components effectively with a single fan, even going at full speed. And the single fan in the power supply was basically designed to cool the power supply, not the components in the case. This is why processors and some graphics cards have their own fans now.

A few years ago PCs often had two or more case fans in addition to the power-supply fan. This made sense in the era of 80mm fans and lots of expansion cards obstructing the airflow, but it was noisy. Nowadays, with sound and graphics and Ethernet integrated onto motherboards, expansion cards are much less common (and processors carry their own mini-fans). Thus, today's standard is to mount one 120mm fan, usually low and forward just beneath the disk-drive stack. This is much quieter, like by a factor of three or four.

The noise produced by a fan is not just a function of the speed with which it turns. It also depends on the nature of the airflow produced by the fan blades and the bearings of the rotor. If the blades cause lots of turbulent airflow, the fan produces lots of noise. One brand of fans that is much more silent than most others even if going at full throttle is Papst [<http://www.papstplc.com/>].

Motherboards

Provided you exercise a little prudence and stay out of the price basement, motherboards and BIOS chips don't vary much in quality. There are only six or so major brands of motherboard inside all those cases and they're pretty much interchangeable; brand premiums are low to nonexistent and cost is strictly tied to maximum speed and bus type. There are only four major brands of BIOS chip (AMI, Phoenix, Mylex, Award) and not much to choose between 'em but the look of the self-test screens (even the "name" vendors use lightly customized versions of these). One advantage Unix buyers have is that Unixes are built not to rely on the BIOS code (because it can't be used in protected mode without more pain than than it's worth). If your BIOS will boot properly, you're usually going to be OK.

Some good features to look for in a motherboard include:

- Gold-plated contacts in the expansion slots and RAM sockets. Base-metal contacts tend to grow an oxidation layer which can cause intermittent connection faults that look like bad RAM chips or boards. (This is why, if your hardware starts flaking out, one of the first things to do is jiggle or remove the boards and reseal them, and press down on the RAM chips to reseal them as well —this may break up the oxidation layer. If this doesn't work, rubbing what contacts you can reach with a soft eraser is a good fast way to remove the oxidation film. Beware, some hard erasers, including many pencil erasers, can strip off the plating, too!)
- The board should be speed-rated as high as your processor, of course. It's good if it's rated higher, so upgrade to a faster processor is just a matter of dropping in the chip and a new crystal.

(I used to have "Voltage, temperature and fan speed monitoring hardware." on this list. But processors run so hot nowadays that all current motherboards have it.)

The dominant form factor is still ATX. Intel tried to replace it with a new standard called BTX in late 2004-2005, but failed; the proposal was effectively withdrawn in 2006. In January 2007 AMD announced a DTX [<http://www.dtxpc.org/>] specification for small-form-factor PCs; it seems also to have sunk without trace.

Monitor and Video

The largest user-visible change since the last major update of this guide is that the CRT (cathode-ray tube) is dead. The manufacturers shut down their production lines in late 2004; the remaining CRTs out there are old stock that's been sitting in warehouses. The only reason to buy one since then has been to get high-end

resolution at a price lower than the insanely expensive high-end flatscreens; with 1920x1440 flatscreens having become generally available at reasonable prices even that reason is gone. It's all flatscreens now, baby.

On flatscreens, only two statistics matter; pixel size and response time. The biggest functional drawback of flatscreens relative to CRTs is that they refresh more slowly, because chemical reactions in a flatscreen pixel take longer than remodulating a flying electron beam. You'll never notice this during ordinary desktop use, but it can cause streakiness and artifacts when you're playing games or viewing movies. If you're going to do that a lot, the price premium for a flatscreen with better response time may be worth it.

Next, buy your card (if you have to; see next paragraph). This used to be complicated, with issues like matching the video bandwidths of the card and the CRT, and the amount of display memory. Now (unless you are a gamer or have similarly extreme 3D acceleration requirements) it's simple; all cards have enough display memory for every resolution in use, and the issues are software (does it have an open-source driver, and do you care?)

It's actually fairly likely you'll never buy a video card again. Very capable graphics chips are routinely integrated onto motherboards now; unless you're a gamer or somebody else who absolutely must have the latest wheeze in 3D acceleration, they'll be good enough. Even this is not a stable situation, as 3D acceleration is commoditizing too.

I used to carry a lot of material on different video standards, interlacing, and flicker. That stuff is all obsolete now.

Here's what to look for on the monitor spec sheet:

- Screen size and format. Usually measured in diagonal inches. Most displays are now in a "widescreen" format (16:10 ratio of width:height) rather than the older 5:4 or 4:3 ratios common for CRTs and older flat panel screens. A "19 inch" widescreen monitor generally has considerably fewer pixels than a "19 inch" 5:4 ratio one. Unfortunately, this change is bad for programmers, as it tends to lose us the vertical pixel resolution we want for editor windows.
- Screen resolution. 1280x1024 is now low end on the desktop. Seventeen inch 1280x1024 screens are the bargain basement now, many manufacturers have already switched production to 19 inch widescreen 1440x900 screens instead. The cost difference between such screens and 20 inch 1680x1050 screens is very small, making the 20 inch screens a better choice unless funds (or desktop space!) are very tight.
- 5ms or lower response time. 3ms is better. There is some marketing-speak going on in the way the response time is specified (grey to grey rather than black to white) but since most manufacturers do it this way these times are usually comparable between different manufacturers screens.
- Does it have a tilt-and-swivel base? Adequate controls, including both horizontal and vertical size and horizontal and vertical centering? A color-temperature control is a plus; the last is particularly important if you compose graphics on screen for hardcopy from a printer.

If you can, buy your monitor from someplace that will let you see the same monitor (the very unit you will walk out the door with, not a different or 'demo' unit of the same model) that will be on your system. There's significant quality variation (even in "premium" monitor brands) even among monitors of the same make and model.

DVD Drives

DVD drives have two main uses in computer systems: playback of video DVDs, and use for data storage (either installation media or backups, or even as a primary drive in a few specialized systems).

DVD video playback used to be problematic on Unix due to various stupid copy-protection schemes in firmware, but they have long since been cracked. These days, any SATA DVD will do fine.

DVD burners (drives that can read and write CDROM media as well as several kinds of DVD media) are now low cost and useful. The SATA interface has taken over here, too. Linux and most current PC Unix-like systems will work fine with either interface, which is good as most PCs now ship with one.

Sound Cards and Speakers

You can't buy a really bad sound card any more. Even low-end sound cards or the sound chips embedded in a lot of PC motherboards these days support all these features:

- 16-bit sampling (for 65536 dynamic levels rather than 256).
- Mono and stereo support.
- Full-duplex mode.
- Sampling rate of 44.1KHz (CD-quality).
- MIDI interface via a standard 15-pin D-shell connector.
- RCA output jacks for headphones or speakers.
- A microphone jack for sound input.

If you are interested in multi-track digital recording, two particularly good choices are the M-Audio Delta, or RME Hammerfall series of cards. Decent (and lower cost!) two-channel cards for more normal use are those using the ICE1712 (Envy24) and ICE1724 (Envy24HT) audio chips. For normal users, though, the on-motherboard chips will work fine.

A rather comprehensive list of sound cards and chips supported by the ALSA project, which is the main way sound cards are supported under Linux, can be found at ALSA Sound Card Matrix [<http://www.alsa-project.org/main/index.php/Matrix:Main>].

In speakers, look for a magnetically-shielded enclosure with volume, bass and treble controls. Some speakers run off the card's 4-watt signal; others are "self-powered", using batteries or a separate power supply. Your major buying choice is which one of these options to pursue. Usually you'll want separately-powered speakers. If appropriate for your listening habits, a pair of decent headphones will get you better quality sound for the money compared to speakers.

One final, important tip: that audio cable from your CD-ROM back to the sound card is used only when you play audio CD-ROMs through your speakers. Software-generated sound goes through the system bus, so you can play games with sound even if your sound board or motherboard won't accept the audio cable connector.

Modems

Demand for (dialup telephone) modems is dropping as more and more people get broadband Internet through DSL and cable. This section still has as much detail as it does only because (a) there are people out beyond the exurbs who can't get broadband, and (b) there are one or two remaining traps for the unwary.

The modem market has stabilized and standardized. If you can spend \$59, get a U.S. Robotics V.92 USB external. You can then know that you've got the best and skip the rest of this section. If you really must

economize, spend \$39 for the internal-card version (but you'll probably regret the \$20 first time you have to do diagnostics).

Note

If you live somewhere with *really* bad telephone lines, the U.S. Robotics V.92 Business Modem may be truly "the best" for your needs, though it is about four times the price of the U.S. Robotics V.92 USB external, which is marketed for home use. See the U.S. Robotics [<http://www.usrobotics.com>] web site for current product numbers and more detailed specifications.

The modem market is like consumer electronics (and unlike the computer market as a whole) in that price is a very poor predictor of performance. For ordinary file transfers, some \$50 modems are better than some \$150 modems. Paying top dollar mainly buys you better tolerance of poor connections and better performance at heavy-duty bi-directional transfers (such as you would generate, for example, using SLIP or PPP over a leased line to an Internet provider).

In today's market all modems do a nominal 56kbps — V.90 and V.92 plus V.29 or V.17 fax transmission and reception (over plain old phone lines you won't get more than 53K of that). You don't see much in the way of slow/cheap to fast/expensive product ranges within a single brand, because competition is fierce and for many modem board designs (those featuring DSP (Digital Signal Processor) chips run by a program in ROM) adding a new protocol is basically a software change.

Internal vs. External

Most modems come in two packagings: internal, designed to fit in a PC card slot, and external, with its own case, power supply, and front-panel lights. Typically you'll pay \$20 to \$30 more for an external modem than you will for the internal equivalent. You'll also need a serial or USB port to connect your external modem to.

Pay that premium — being able to see the blinkenlights on the external ones will help you understand and recover from pathological situations. For example, if your Unix system is prone to "screaming-tty" syndrome, you'll quickly learn to recognize the pattern of flickers that goes with it. Punch the hangup/reset button on an external modem and you're done — whereas with an internal modem, you have to go root and flounder around killing processes and maybe cold-boot the machine just to reset the card.

See Rick's Rants [<http://linuxmafia.com/~rick/faq/>] for extended discussion of this point.

Pitfalls to Avoid

Don't buy a serial (RS232C) modem. This used to be the only kind there was, but they were always a bitch to configure and troubleshoot. Go USB instead; the sanity you save may be your own.

If the abbreviation "RPI" occurs anywhere on the box, don't even consider buying the modem. RPI (Rockwell Protocol Interface) is a proprietary "standard" that allows modem makers to save a few bucks at your expense by using a cheap-jack Rockwell chipset that doesn't do error correction. Instead, it hands the job off to a modem driver which (on a Unix machine) you will not have.

Also avoid anything called a "Windows Modem" or "WinModem", "HCF", or "HSP"; these lobotomized pieces of crap require a Windows DLL to run. They will eat up to 25% of your processor clocks during transfers, and hog high-priority interrupts (causing your machine to stall under Windows even if your processor still has spare cycles).

A good way to avoid falling into the WinModem trap is to look for the designation "OEM modem". This is apparently the new industry-speak for a modem with an on-board hardware DSP. Occasionally you'll see these called "gaming modems".

Fax Modems

Many modems come with bundled Windows fax software that is at best useless under Unix, and at worst a software kluge to cover inadequate hardware. Avoid these bundles and buy a bare modem — it's cheaper, and lowers the likelihood that something vital to your communications needs has been left out of the hardware.

Avoid “Class 1” and “Class 2” modems. Look for “Class 2.0” for the full EIA-standard command set.

Fax capability is included with effectively all modems these days; it's cheap for manufacturers, being basically a pure software add-on. The CCITT also sets fax protocol standards. Terms to know:

V.29 CCITT standard for Group III fax encoding at 9600bps

V.17 CCITT standard for Group III fax encoding at 14400bps

There's a separate series of standards for software control of fax modems over the serial (or USB) line maintained by the Electronics Industry Association and friends. These are:

Class 1 — base EIA standard for fax control as extensions to the Hayes AT command set.

Class 2.0 — enhanced EIA standard including compression, error correction, station ID and other features.

Class 2 — marketroidian for anything between Class 1 and Class 2.0. Different “Class 2” modems implement different draft subsets of the 2.0 standard, so “Class 2” fax software won't necessarily drive any given “Class 2” modem.

There's also a proprietary Intel “standard” called CAS, Communicating Applications Specification. Ignore it; only Intel products support it.

Printers

The most important thing to optimize nowadays is cost of consumables. Printer manufacturers, especially at the low end, have adopted a model under which they sell printers with near-zero or even negative margin, then gouge you horribly on the cost of cartridges and ink. Common tactics include (a) shipping half-filled “starter” cartridges with your printer, so you have to replace much sooner than you'd think, (b) toner-empty sensors deliberately miscalibrated to blink the error light on your printer when they're still a quarter to a third full, and (c) electronic countermeasures to lock out cheap third-party refills - in one notorious case, a printer manufacturer used the DMCA to sue refill vendors who circumvented these!

Better dealers (the Staples chain, for example) will show you a chart covering price and consumable-cost-per-page for all the models they carry. If you don't see this, leave. When you do, estimate your monthly print volume and trade off up-front against consumables price. appropriately. Hint: The vendors count on you underestimating your volume and consumables cost, and you usually will. Paying a few extra bucks up front to lower that cost is smart.

Other than that, there really isn't all that much to be said about printers; the market is thoroughly commoditized and printer capabilities pretty much independent of the rest of your hardware. The PC-clone magazines will tell you what you need to know about print quality, speed, features, etc. The business users they feed on are obsessed with all these things.

(There used to be a problem with “GDI printers” and “WinPrinters” that only worked with Windows — they required special drivers that took over your CPU to do image processing. These were such a bad idea that they have basically disappeared off the market.)

Most popular printers are supported by GhostScript, and so it's easy to make them do PostScript. If you're buying any letter-quality printer (laser or ink-jet), check to see if it's on GhostScript's supported device list — otherwise you'll have to pay a premium for Postscript capability! Postscript is still high-end in the Windows market, but it's ubiquitous in the Unix world.

Warning, however: if you're using ghostscript on a non-Postscript printer, printspeed will be slow, especially with a serial printer. A bitmapped 600 dpi page has a *lot* of pixels on it. At today's prices, paying the small premium for Postscript capability makes sense.

If you're buying a printer for home, an inkjet is a good choice because it doesn't use gobs of power and you won't have the toner/ozone/noise/etc mess that you do with a laser. If all you want is plain-ASCII, dot-matrix is cheaper to buy and run — if you can find one. Inexpensive ink-jets and lasers have almost driven them off the market.

Inkjets are great in that they're cheap, many of them do color, and there are many kinds which aren't PCL but are understood by Ghostscript anyway. If you print very infrequently (less than weekly, say), you should be careful to buy a printer whose print head gets replaced with every ink cartridge: infrequent use can lead to the drying of the ink, both in the ink cartridge and in the print head. The print heads you don't replace with the cartridge tend to cost nearly as much as the printer (~\$200 for an Epson Stylus 800) once the warranty runs out (the third such repair, just after the warranty expired, totalled one informant's Stylus 800). Be careful, check print head replacement costs ahead of time, and run at least a cleaning cycle if you don't actually print anything in a given week. (Conversely, toner starts out dry, and ribbon ink won't evaporate for years...if you truly print only rarely, but neither a dot matrix nor a laser makes sense, consider buying no printer and taking your PostScript files to a copy shop...)

Nowadays, a lot of printers are moving away from parallel-port interfaces to USB. This is a good idea, because USB devices announce themselves to the host computer and can be automatically configured. Parallel ports (and serial ports for that matter) are becoming obsolete. Many new PC motherboards no longer include them.

Many printers (even some sub-\$100 models) now come with a network (10/100 Ethernet) interface. This makes sharing them trivial, and also avoids having to leave a desktop PC powered on so others (using notebooks perhaps) can print to your printer. Therefore, such printers are worth considering in many networked environments, including home networks.

In the near future, new motherboards may stop including parallel and serial ports altogether. That's another good reason to go with a USB- or Ethernet-capable printer.

Power Protection

I strongly recommend that you buy a UPS to protect your hardware and data. MOV-filtered power bars make nice fuses (they're cheap to replace), but they're not enough. I have written a UPS HOWTO [<http://www.tldp.org/HOWTO/UPS-HOWTO>] that provides more complete coverage of what used to be in this section.

Radio Frequency Interference

(Thanks to Robert Corbett <Robert.Corbett@Eng.Sun.COM> for contributing much of this section)

Radio Frequency Interference (RFI) is a growing problem with PC-class machines. Today's processor speeds are such that the electromagnetic noise generated by a PC's circuitry in normal operation can degrade or jam radio and TV reception in the neighborhood. Such noise is called Radio Frequency Interference (RFI). Computers, as transmitting devices, are regulated by the Federal Communications Commission (FCC).

FCC regulations recognize two classes of computer:

If a PC is to be used in a home or apartment, it must be certified to be FCC class B. If it is not, neighbors have a legal right to prevent its use. FCC class A equipment is allowed in industrial environments.

Many systems are not FCC class B. Some manufacturers build boxes that are class B and then ship them with class A monitors or external disk drives. Even the cables can be a source of RFI.

It pays to be cautious. For example, the Mag MX17F is FCC class B. There are less expensive versions of the MX17 that are not. The Mag MX17 is a great monitor. It would be painful to own one and not be allowed to use it.

An upgradeable system poses special problems. A system that is FCC class B with a 33 MHz CPU might not be when the CPU is upgraded to a 50 or 66 MHz CPU. Some upgrades require knockouts in the case to be removed. If a knockout is larger than whatever replaces it, RFI can leak out through the gap. Grounded metal shims can eliminate the leaks.

Even Class B systems don't mix well with wireless phonesets (not cellular phones, but the kind with a base station and antennaed headset). You'll often find a wireless phone hard to use withing 20 feet of a Class B machine.

To cut down on RFI, get a good metal case with tight joints, or at least make sure any plastic one you buy has a conductive lining. You can also strip the painted metal-to-metal contacting parts of paint so that there's good conductive metal contact. Paint's a poor conductor in most cases, so you can get some benefit from this.

What To Optimize

First, add more memory

Max out your memory. Having lots of free memory will improve your virtual-memory performance (and Unix takes advantage of extra memory more effectively than Windows does). Fortunately, with RAM as cheap as it is now, a gigabyte or three is unlikely to bust your budget even if you're economizing.

Bus and Disk speeds

Most people think of the processor as the most important choice in specifying any kind of personal-computer system. But for typical job loads under Linux, the processor type is nearly a red herring — it's far more important to specify a capable system bus and disk I/O subsystem. If you don't believe this, you may find it enlightening to keep `top(1)` running for a while as you use your machine. Notice how seldom the CPU idle percentage drops below 90%!

It's true that after people upgrade their motherboards they often do report a throughput increase. But this is often more due to other changes that go with the processor upgrade, such as improved cache memory or an increase in the clocking speed of the system's front-side bus (enabling data to get in and out of the processor faster).

If you're buying for Linux on a fixed budget, it makes sense to trade away some excess processor clocks to get a faster bus and disk subsystem. If you're building a monster hot-rod, go ahead and buy that fastest available processor — but once you've gotten past that gearhead desire for big numbers, pay careful attention to bus speeds and your disk subsystem, because that's where you'll get the serious performance wins. The gap between processor speed and I/O subsystem throughput has only widened in the last five years.

How does it translate into a recipe in 2010? Like this; if you're building a hot rod,

- *Do* buy a machine with the fastest available "front-side" (e.g. processor-to-memory) bus.
- *Do* get the fastest SATA disks you can get your hands on.

Optimizing your disk subsystem

For the fastest disks you can find, pay close attention to average seek and latency time. The former is an average time required to seek to any track; the latter is the maximum time required for any sector on a track to come under the heads, and is a function of the disk's rotation speed.

Of these, average seek time is *much* more important. When you're running Linux or any other virtual-memory operating system, a one millisecond faster seek time can make a really substantial difference in system throughput. Back when PC processors were slow enough for the comparison to be possible (and I was running System V Unix), it was easily worth as much as a 30MHz increment in processor speed. Today the corresponding figure would probably be as much as 300MHz!

The manufacturers themselves avoid running up seek time on the larger-capacity drives by stacking platters vertically rather than increasing the platter size. Thus, seek time (which is proportional to the platter radius and head-motion speed) tends to be constant across different capacities in the same product line. This is good because it means you don't have to worry about a capacity-vs.-speed tradeoff.

Average drive latency is inversely proportional to the disk's rotational speed. For years, most disks spun at 3600 rpm; most disks now spin at 7,200 or 10,000rpm, and high-end disks at 15,000 rpm. These fast-spin disks run extremely hot; cooling is becoming a critical constraint in drive design.

For years, your basic decision was SATA vs. SCSI (the older IDE and EIDE buses are long obsolete). Not in 2009; SATA 3 devices and controllers are good enough that the performance advantage of SCSI is marginal unless you are designing a super-high-end server box - slightly faster transfer speeds (320MB/s vs. 300MB/s) and slightly better sustained throughput.

The SCSI price premium entailed in an extra controller and more expensive disks are no longer worth it for the home builder, even from the point of view of grizzled old SCSI partisans like me. Accordingly, I've dropped most of the detailed SCSI information I used to carry here.

Final note: Solid-state drives loom on the horizon as replacements for SATA disks, but the price per megabyte is still high enough that as yet they're only being deployed in small capacities on netbooks. Watch this space.

Tuning Your I/O Subsystem

(This section comes to us courtesy of Perry The Cynic, <perry@str.cynic.org>; it was written in 1998. My own experience agrees pretty completely with his. I have revised the numbers in it since to reflect more recent developments.)

Building a good I/O subsystem boils down to two major points: *pick matched components* so you don't over-build any piece without benefit, and *construct the whole pipe such that it can feed what your OS/application combo needs*.

It's important to recognize that "balance" is with respect to not only a particular processor/memory subsystem, but also to a particular OS and application mix. A Unix server machine running the whole TCP/IP server suite has radically different I/O requirements than a video-editing workstation. For the "big boys" a good consultant will sample the I/O mix (by reading existing system performance logs or taking new measurements) and figure out how big the I/O system needs to be to satisfy that app mix. This is not something your typical Linux buyer will want to do; for one, the application mix is not static and will

change over time. So what you'll do instead is design an I/O subsystem that is internally matched and provides maximum potential I/O performance for the money you're willing to spend. Then you look at the price points and compare them with those for the memory subsystem. That's the most important trade-off inside the box.

So the job now is to design and buy an I/O subsystem that is well matched to provide the best bang for your buck. The two major performance numbers for disk I/O are latency and bandwidth. Latency is how long a program has to wait to get a little piece of random data it asked for. Bandwidth is how much contiguous data can be sent to/from the disk once you've done the first piece. Latency is measured in milliseconds (ms); bandwidth in megabytes per second (MB/s). Obviously, a third number of interest is how big all of your disks are together (how much storage you've got), in Gigabytes (GB).

Within a rather big envelope, minimizing latency is the cat's meow. Every millisecond you shave off effective latency will make your system feel significantly faster. Bandwidth, on the other hand, only helps you if you suck a big chunk of contiguous data off the disk, which happens rarely to most programs. You have to keep bandwidth in mind to avoid mis-matching pieces, because (obviously) the lowest usable bandwidth in a pipe constrains everything else.

I'm going to ignore IDE. IDE is no good for multi-processing systems, period. You may use an IDE CD-ROM if you don't care about its performance, but if you care about your I/O performance, go SCSI. (Beware that if you mix an IDE CD-ROM with SCSI drives under Linux, you'll have to run a SCSI emulation layer that is a bit flaky.)

Let's look at the disks first. Whenever you seriously look at a disk, *get its data sheet*. Every reputable manufacturer has them on their website; just read off the product code and follow the bouncing lights. Beware of numbers (<12ms fast!) you may see in ads; these folks often look for the lowest/highest numbers on the data sheet and stick them into the ad copy. Not dishonest (usually), but ignorant.

What you need to find out for a disk is:

1. What kind of SCSI interface does it have? Look for "fast", "ultra", and "wide". Ignore disks that say "fiber" (this is a specialty physical layer not appropriate for the insides of small computers). Note that you'll often find the same disk with different interfaces.
2. What is the "typical seek" time (ms)? Make sure you get "typical", not "track-to-track" or "maximum" or some other measure (these don't relate in obvious ways, due to things like head-settling time).
3. What is the rotational speed? This is typically 4500, 5400, 7200, 10000, or 15000 rpm (rotations per minute). Also look for "rotational latency" (in ms). (In a pinch, average rotational latency is approx. 30000/rpm in milliseconds.)
4. What is the 'media transfer rate' or speed (in MB/s)? Many disks will have a range of numbers (say, 7.2-10.8MB/s). Don't confuse this with the "interface transfer rate" which is always a round number (10 or 20 or 40MB/s) and is the speed of the SCSI bus itself.

These numbers will let you do apple-with-apples comparisons of disks. Beware that they will differ on different-size models of the same disk; typically, bigger disks have slower seek times.

Now what does it all mean? Bandwidth first: the 'media transfer rate' is how much data you can, under ideal conditions, get off the disk per second. This is a function mostly of rotation speed; the faster the disk rotates, the more data passes under the heads per time unit. This constrains the sustained bandwidth of *this disk*.

More interestingly, your effective latency is the sum of typical seek time and rotational latency. So for a disk with 8.5ms seek time and 4ms rotational latency, you can expect to spend about 12.5ms between the moment the disk 'wants' to read your data and the moment when it actually starts reading it. This is

the one number you are trying to make small. Thus, you're looking for a disk with low seek times and high rotation (RPM) rates.

For comparison purposes, the first hard drive I ever bought was a 20MB drive with 65ms seek time and about 3000RPM rotation. A floppy drive has about 100-200ms seek time. A CD-ROM drive can be anywhere between 120ms (fast) and 400ms (slow). The best IDE harddrives have about 10-12ms and 5400 rpm. The best SCSI harddrive I know (the Fujitsu MAM) runs 3.9ms/15000rpm.

Fast, big drives are expensive. Really big drives are very expensive. Really fast drives are pretty expensive. On the other end, really slow, small drives are cheap but not cost effective, because it doesn't cost any less to make the cases, ship the drives, and sell them.

In between is a 'sweet spot' where moving in either direction (cheaper or more expensive) will cost you more than you get out of it. The sweet spot moves (towards better value) with time. If you can make the effort, go to your local computer superstore and write down a dozen or so drives they sell 'naked'. (If they don't sell at least a dozen hard drives naked, find yourself a better store. Use the Web, Luke!) Plot cost against size, seek and rotational speed, and it will usually become pretty obvious which ones to get for your budget.

Do look for specials in stores; many superstores buy overstock from manufacturers. If this is near the sweet spot, it's often surprisingly cheaper than comparable drives. Just make sure you understand the warranty procedures.

Note that if you need a lot of capacity, you may be better off with two (or more) drives than a single, bigger one. Not only can it be cheaper but you end up with two separate head assemblies that move independently, which can cut down on latency quite a bit (see below).

If you find yourself at the high end of the bandwidth game, be aware that the theoretical maximum of the PCI bus itself is 132MB/s. That means that a dual ultra/wide SCSI controller (2x40MB/s) can fill more than half of the PCI bus's bandwidth, and it is not advised to add another fast controller to that mix. As it is, your device driver better be well written, or your entire system will melt down (figuratively speaking).

Incidentally, all of the numbers I used are 'optimal' bandwidth numbers. The real scoop is usually somewhere between 50-70% of nominal, but things tend to cancel out — the drives don't quite transfer as fast as they might, but the SCSI bus has overhead too, as does the controller card.

Whether you have a single disk or multiple ones, on one or several SCSI buses, you should give careful thought to their partition layout. Given a set of disks and controllers, this is the most crucial performance decision you'll make.

A partition is a contiguous group of sectors on the disk. Partitioning typically starts at the outside and proceeds inwards. All partitions on one disk share a single head assembly. That means that if you try to overlap I/O on the first and last partition of a disk, the heads must move full stroke back and forth over the disk, which can radically increase seek time delays. A partition that is in the middle of a partition stack is likely to have best seek performance, since at worst the heads only have to move half-way to get there (and they're likely to be around the area anyway).

Whenever possible, split partitions that compete onto different disks. For example, /usr and the swap should be on different disks if at all possible (unless you have outrageous amounts of RAM).

Another wrinkle is that most modern disks use 'zone sectoring'. The upshot is that outside partitions will have higher bandwidth than inner ones (there is more data under the heads per revolution). So if you need a work area for data streaming (say, a CD-R pre-image to record), it should go on an outside (early numbered) partition of a fast-rotating disk. Conversely, it's a good convention to put rarely-used, performance-uncritical partitions on the inside (last).

Ah yes, caches. There are three major points where you could cache I/O buffers: the OS, the controller card or chip in your machine, and the on-disk controller. Intelligent OS caching is by far the biggest win, for many reasons. RAM caches on controller cards are pretty pointless these days; you shouldn't pay extra for them, and experiment with disabling them if you're into tinkering.

RAM caches on the drives themselves are a mixed bag. At moderate size (1-2MB), they are a potential big win for Windows 95/98, because Windows has stupid VM and I/O drivers. If you run a true multi-tasking OS like Linux, having unified RAM caches on the disk is a significant loss, since the overlapping I/O threads kick each other out of the cache, and the disk ends up performing work for nothing.

Most high-performance disks can be reconfigured (using mode page parameters, see above) to have 'segmented' caches (sort of like a set-associative memory cache). With that configured properly, the RAM caches can be a moderate win, not because caching is so great on the disk (it's much better in the OS), but because it allows the disk controller more flexibility to reschedule its I/O request queue. You won't really notice it unless you routinely have >2 I/O requests pending at the SCSI level. The conventional wisdom (try it both ways) applies.

And finally I *do* have to make a disclaimer. Much of the stuff above is shameless simplification. In reality, high-performance disks are very complicated beasts. They run little mini-operating systems that are most comfortable if they have 10-20 I/O requests pending *at the same time*. Under those circumstances, the amortized global latencies are much reduced, though any single request may experience *longer* latencies than if it were the only one pending. The only really valid analysis are stochastic-process models, which we *really* don't want to get into here. :-)

But What If I'm Economizing?

If you are economizing, here's a simple rule:

- *Do* buy a CPU/motherboard one or two levels lower than commercial state of the art.

For best value, look in the *middle* of the current range of available processors. On the desktop, in late 2007, that means a CPU costing perhaps \$75 to \$200, not the latest and greatest quad core marvels selling for several times that!

Why? Because of the way manufacturers' price-performance curves are shaped. The top-of-line system is generally boob bait for corporate executives and other people with more money than sense. Chances are the system design is new and untried — if you're at the wrong point in the technology cycle, the chip may even be a pre-production sample, or an early production stepping with undiscovered bugs like the infamous FDIV problem in early Pentiums. You don't need such troubles. Better to go with a chip/motherboard combination that's been out for a while and is known good. It's not like you really need the extra speed, after all.

Besides, if you buy one of these gold-plated systems, you're only going to kick yourself three months later when the price plunges by 30%. Further down the product line there's been more real competition and the manufacturer's margins are already squeezed. There's less room for prices to fall, so you won't watch your new toy lose street value so fast. Its price will still drop, but it won't plummet sickeningly.

Again, bear in mind that the cheapest processor you can buy new today is plenty fast enough for Linux. So if dropping back a speed level or two brings you in under budget, you can do it with no regrets.

Consider one drive rather than two. This *will* reduce overall system performance somewhat, but the cost saving as a fraction of total system cost is often substantial.

Another easy economy measure is looking for repaired or reconditioned parts with a warranty. These are often as good as new, and much cheaper.

Your display is one of the areas where pinching pennies is *not* a good idea. You're going to be looking at that display for hours on end. You are going to be using the screen real estate constantly. Buy the best quality, largest screen you possibly can — it will be worth it.

Similarly, do not reduce the amount of RAM in your system too far. A minimum of 4GB of RAM is helpful in desktop systems today.

Noise Control and Heat Dissipation

An increasingly critical aspect of machine design is handling the waste heat and acoustic noise of operation. This may seem like a boring subject, but cooling is a centrally important one if you want your machine to last — because thermal stress from the electronics' own waste heat is almost certainly what will kill it. You want that fatal moment to happen later rather than sooner. On the other hand, cooling makes acoustic noise, which human beings don't tolerate well.

This tradeoff bites harder than you think; it's the fundamental reason that, despite my money-is-no-object premise in the Ultimate Linux Box articles, I didn't go to relatively exotic technologies like liquid-cooled overclocking or RAID disk arrays for a performance boost. Sure, they may initially look attractive — but overclocked chips and banks of disk drives require massive cooling with lots of moving parts, and it's not good to be trying to do creative work like programming with anything that sounds quite so much like an idling jet engine sitting beside one's desk.

In 2001 we had already reached the point at which the thermal load vs. cooling-noise tradeoff is the effective limiting factor in the performance of personal machines. Ten years ago, even low-end and medium "server" machines differed from personal-PC designs in fairly important ways (different processor and bus types, different speed ranges, etc.) Nowadays specialized server architectures are in retreat at the high end of the market and everything else looks like a PC. And the difference between a "PC" and a "server" is mainly that servers live in server rooms, and are allowed to have monster cases with lots of noisy fans.

So how do we manage this tradeoff for a personal, desktop or desk-side machine? Careful choice of components and being willing to pay some price premium for cool-running and low-noise characteristics can help a lot. Even exceptionally clueful system integrators can't generally afford to do this, because they're under constant competitive pressure to cut price and costs by using generic components.

Reducing expected noise and heat in a design call for different strategies. It's relatively easy to find decibel figures for the noisemaking parts in a PC design. And, once you know a little basic audiometry and a few basic rules of thumb, it's not hard to form a fair estimate of your design's noisiness. Estimating a design's heat dissipation is harder, partly because the waste-heat emission of a PC's subsystems tends to vary in a more complex way than the acoustic emissions of the mechanical parts. This means that you can and should try to design ahead for low noise, but on the other hand expect to have to monitor for heat-dissipation problems in your prototype and solve them by building in more cooling.

Here's the basic audiometry you need to know to control your design's noise emissions:

Sound is measured in *decibels*, abbreviated dB, relative to the threshold of audibility, "A". (Thus, sound levels above that threshold are written "dBA".) The scale is logarithmic, with every 3dB increment roughly doubling sound intensity.

For sounds that are not phase-related, decibel levels add as a logarithmic sum. Thus if X and Y are uncorrelated sound sources,

$$\text{dBA}(X + Y) = 10 * \log(10 ^ { (\text{dBA}(X)/10) } + 10 ^ { (\text{dBA}(Y)/10) })$$

A consequence of the above formula is that $\text{dBA}(X + Y)$ cannot be more than 3dB above the greater of $\text{dBA}(X)$ and $\text{dBA}(Y)$ for uncorrelated sources (6dB for perfectly correlated ones).

Sound from a point source decays by an inverse-square law, roughly 6dB for each doubling of distance.

Important thresholds on the decibel scale:

0 dBA	Threshold of hearing
20 dBA	Rustling leaves, quiet living room
30 dBA	Quiet office
40 dBA	Quiet conversation
45 dBA	Threshold of distraction, according to EPA
50 dBA	Quiet street, average office noise
60 dBA	Normal conversation (1 foot distance)
70 dBA	Inside car
75 dBA	Loud singing (3 feet)
80 dBA	Typical home-stereo listening level

The acoustic noise emitted by PCs is normally a combination of white noise produced by airflow, high-frequency noise produced by bearing friction in drive spindles and fans, and the constant frequency "blade passing" noise that all propellers emit (the latter is often more intense than white noise and bearing whine).

The best low-noise ball-bearing case fans emit around 20dBA. Typical sleeve-bearing fans emit 30-50dBA.

According to the indispensable Tom's Hardware site [<http://tomshardware.com/>], you can expect to cut at least 5dB off the interior noise level of the computer with a good choice of case. We'll improve on that by adding sound-absorbing material to the interior.

Special Considerations When Buying Laptops and Netbooks

First, don't be misled by the term "netbook". A netbook is just a small, low-priced, low-power laptop with relatively small solid-state drives. Because the display and drive capacity are small, netbooks are basically just good for email and surfing. If you're going to do coding or even much word processing you'll need something more like a traditional laptop or desktop.

Up until about 1999 the laptop market was completely crazy. The technology was in a state of violent flux, with "standards" phasing in and out and prices dropping like rocks. Things are beginning to settle out a bit more now.

One sign of this change is that there are now a couple of laptop lines that are clear best-of-breeds for reasons having as much to do with good industrial design and ergonomics as the technical details of the processor and display.

In lightweight machines, I was a big fan of the Sony VAIO line. I owned one from early 1999 until it physically disintegrated under the rigors of travel in late 2000, and could hardly imagine switching. They weigh 3.5 pounds, give you an honest 3 hours of life per detachable battery pack, have a very nice 1024x768 display, and are just plain *pretty*. Their only serious drawback is that they're not rugged, and often fall apart after a year or so of use.

If you want a full-power laptop that can compete with or replace your desktop machine, the Lenovo (formerly IBM) ThinkPad line is the bomb. Capable, rugged, and nicely designed. I now use a ThinkPad X61, the lightest and smallest machine in the line, and like it a lot.

These machines are not cheap, though. If you're trying to save money by buying a no-name laptop, here are things to look for:

First: despite what you may believe, the most important aspect of any laptop is *not* the CPU, or the disk, or the memory, or the screen, or the battery capacity. It's the keyboard feel, since unlike in a PC, you cannot throw the keyboard away and replace it with another one unless you replace the whole computer. *Never buy any laptop that you have not typed on for a couple hours.* Trying a keyboard for a few minutes is not enough. Keyboards have very subtle properties that can still affect whether they mess up your wrists.

A standard desktop keyboard has keycaps 19mm across with 7.55mm between them. If you plot frequency of typing errors against keycap size, it turns out there's a sharp knee in the curve at 17.8 millimeters. Beware of "kneetop" and "palmtop" machines, which squeeze the keycaps a lot tighter and typically don't have enough oomph for Unix anyway; you're best off with the "notebook" class machines that have full-sized keys.

Second: with present flatscreens, 1920x1200 color is the best you're going to do (and that is on a 17in widescreen, which translates to a large notebook. On normal size notebooks, a maximum of 1440x900 is more common). On travel machines like the Lenovo X series, you're still stuck with 1024x768. If you want more than that (for X, for example) you have to either fall back to a desktop or make sure there's an external-monitor port on the laptop (and many laptops won't support higher resolution than the flatscreen's).

Third: about those vendor-supplied time-between-recharge figures; *don't believe them.* They collect those from a totally quiescent machine, sometimes with the screen or hard disk turned off. Under Windows, you'd be lucky to get half the endurance they quote; under Unix, which hits the disk more often, it may be less yet. Figures from magazine reviews are more reliable.

Fourth: You can now avoid many of the driver hassles involved in getting some devices on your notebook to work (or week well) under Linux by purchasing a notebook with Linux pre-installed. Dell has recently started to make noise in this regard in the Linux community. Taking this approach limits the set of notebooks you can consider, but the one you get is likely to "just work" (including sound, useful capabilities like suspend/resume, and even hotplugging of external displays and projectors) to a much higher degree under Linux than others.

How to Buy

When to Buy

It used to be that good configurations for Unix were what the market called 'server' machines, with beefed-up I/O subsystems and fast buses. No longer; today's 'servers' are monster boxes with multiple power supplies and processors, gigabytes of memory, and industrial-grade air cooling—they're not really suitable as personal machines. A typical SCSI desktop workstation is as much as you'll need.

Prices keep dropping, so there's a temptation to wait forever to buy. A good way to cope with this is to configure your system on paper, get a couple of initial estimates, then set a trigger price, below the lowest one, at what you're willing to pay. Then watch and wait. When the configuration cost hits your trigger price, place your order.

The advantage of this method is that it requires you to settle in your mind, well in advance, what you're willing to pay for what you're getting. That way, you'll buy at the earliest time you should, and won't stress too much out afterwards as it depreciates.

Before you shop, do your homework. Publications like "Computer Shopper" (and their web site at <http://www.computershopper.com> [<http://www.computershopper.com>]) are invaluable for helping you get a feel for prices and what clonemakers are doing. Another excellent site is ComputerESP [<http://www.computeresp.com>].

Where to Buy

The most important where-to-buy advice is negative. Do *not* go to a traditional, business-oriented storefront dealership. Their overheads are high. So are their prices.

Especially, run —do not walk —away from any outfit that trumpets ‘business solutions’. This is marketing code for the kind of place that will justify a heavy price premium by promising after-sale service and training which, nine times out of ten, will turn out to be nonexistent or incompetent. Sure, they'll give you plush carpeting and a firm handshake from a guy with too many teeth and an expensive watch —but did you really want to pay for that?

There are two major alternatives to storefront dealerships and one minor one. The major ones are mail order and computer superstores. The minor one is computer fairs.

Computer Fairs

I used to be a big fan of hole-in-the-wall stores run by immigrants from the other side of the International Date Line, but most of those places have been driven out of the regular retail game by the superstores and the Web. If you still have one in your neighborhood, you're lucky. I do, as it happens, but that is now unusual; the only place you normally find diaspora Chinese and Indians selling cheap PCs over the counter anymore is at computer fairs. (Usually they're doing it to publicize an Internet/mail-order business.)

You can find good loss-leader deals on individual parts at these fairs (they're especially good places to buy disk drives cheap). But I call them a minor alternative because it's hard to get a custom configuration tuned for Unix built for you at a fair. So you end up, effectively, back in the mail-order or Web channel.

Mail Order

Internet buying makes a lot of sense today for anyone with more technical savvy than J. Random Luser in a suit. Even from no-name vendors, parts and system quality tend to be high and consistent, so conventional dealerships don't really have much more to offer than a warm fuzzy feeling. Furthermore, competition has become so intense that even Internet/mail-order vendors today have to offer not just lower prices than ever before but warranty and support policies of a depth that would have seemed incredible a few years back. For example, many bundle a year of on-site hardware support with their medium- and high-end “business” configurations for a very low premium over the bare hardware.

Note, however, that assembling a system yourself out of parts is *not* likely to save you money over dealing with the Internet/mail-order systems houses. You can't buy parts at the volume they do; the discounts they command are bigger than the premiums reflected in their prices. The lack of any system warranty or support can also be a problem even if you're expert enough to do the integration yourself — because you also assume all the risk of defective parts and integration problems.

Watch out for dealers (Spectrum Trading for one) who charge ridiculous shipping fees. One of our spies reports he bought a hotswappable hard disc drive tray that weighed about 3 lbs. and cost \$250 and they charged \$25 to ship it UPS groud.

Don't forget that (most places) you can avoid sales tax by buying from an out-of-state outfit, and save yourself 6-8% depending on where you live. If you live near a state line, buying from a local outfit you can

often win, quite legally, by having the stuff shipped to a friend or relative just over it. Best of all is a buddy with a state-registered dealer number; these aren't very hard to get and confer not just exemption from sales tax but (often) whopping discounts from the vendors. Hand him a dollar afterwards to make it legal.

(Note: I have been advised that you shouldn't try the latter tactic in Florida —they are notoriously tough on "resale license" holders).

(Note II: The Supreme Court has ruled that states may not tax out-of-state businesses under existing law, but left the way open for Congress to pass enabling legislation. Let's hope the mail-order industry has good lobbyists.)

Computer Superstores

Big chain superstores like CompUSA give you a reasonable alternative to the Web. And there are good reasons to explore it — these stores buy and sell at volumes that allow them to offer prices not far above the Web. (They make back a lot of their margin on computer games and small accessories like mouse pads, cables, and floppy disks.)

Note, however: *Avoid Best Buy*. Horror stories about them are legion — predatory salescritters, incompetent service, routine bait-and-switch tactics.

One thing you should not buy remotely if you can avoid it is a monitor. Monitors are subject to significant quality variations even within the same make and model. Flatscreens have this [problem less than CRTs did, but you don't want a flatscreen with dead pixels. So buy your monitor face-to-face, picking the best out of three or four.

Another good argument for buying at a superstore is that you may have to pay return postage if you ship a system back to the vendor. On a big, heavy system, this can eat your initial price savings.

The only major problem with superstores is that the salespeople who staff them aren't very bright or very clueful (it's a sort of Darwinian reverse-selection effect; these are the guys who are fascinated by computer technology but not smart enough to be techies). Most of them don't know from Linux and are likely to push things like two-button mice that you can't use. Use caution and check your system manifest.

But if you shop carefully and don't fall for one of their name-brand "prestige" systems, you can get prices comparable to Internet/mail-order with the comfort of knowing there's a trouble desk you can drive back to in a pinch. (Also, you *can* see your monitor before you buy!)

Other Buying Tips

You can often get out of paying tax just by paying cash, especially at computer shows. You can always say you're going to ship the equipment out of the state.

A lot of vendors bundle Windows and variable amounts of apps with their hardware. If you tell them to lose all this useless cruft they may shave \$50 or \$100 off the system price.

Questions You Should Always Ask Your Vendor

Minimum Warranty Provisions

The weakest guarantee you should settle for in the mail-order market should include:

- 72-hour burn-in to avoid that sudden infant death syndrome. (Also, try to find out if they do a power-cycling test and how many repeats they do; this stresses the hardware much more than steady burn-in.)
- 30 day money-back guarantee. Watch out for fine print that weakens this with a restocking fee or limits it with exclusions.
- 1 year parts and labor guarantee (some vendors give 2 years).
- 1 year of 800 number tech support (many vendors give lifetime support).

Additionally, many vendors offer a year of on-site service free. You should find out who they contract the service to. Also be sure the free service coverage area includes your site; some unscrupulous vendors weasel their way out with “some locations pay extra”, which translates roughly to “through the nose if you're further away than our parking lot”.

If you're buying store-front, find out what they'll guarantee beyond the above. If the answer is “nothing”, go somewhere else.

Documentation

Ask your potential suppliers what kind and volume of documentation they supply with your hardware. You should get, at minimum, operations manuals for the motherboard and each card or peripheral; also an IRQ list. Skimpiness in this area is a valuable clue that they may be using no-name parts from Upper Baluchistan, which is not necessarily a red flag in itself but should prompt you to ask more questions.

A System Quality Checklist

There are various cost-cutting tactics a vendor can use which bring down the system's overall quality. Here are some good questions to ask:

- If you're buying a factory-configured system, does it have FCC certification? While it's not necessarily the case that a non-certified system is going to spew a lot of radio-frequency interference, certification is legally required — and becoming more important as clock frequencies climb. Lack of that sticker may indicate a fly-by-night vendor, or at least one in danger of being raided and shut down! (For further discussion, see the section on Radio Frequency Interference above.)
- Are the internal cable connectors keyed, so they can't be put in upside down? This doesn't matter if you'll never, ever *ever* need to upgrade or service your system. Otherwise, it's pretty important; and, vendors who fluff this detail may be quietly cutting other corners.

Things to Check when Buying

Tricks and Traps in Warranties

Reading warranties is an art in itself. A few tips:

Beware the deadly modifier “manufacturer's” on a warranty; this means you have to go back to the equipment's original manufacturer in case of problems and can't get satisfaction from the mail-order house. Also, manufacturer's warranties run from the date *they* ship; by the time the mail-order house assembles and ships your system, it may have run out!

Watch for the equally deadly “We do not guarantee compatibility”. This gotcha on a component vendor's ad means you may not be able to return, say, a video card that fails to work with your motherboard.

Another dangerous phrase is “We reserve the right to substitute equivalent items”. This means that instead of getting the high-quality name-brand parts advertised in the configuration you just ordered, you may get those no-name parts from Upper Baluchistan — theoretically equivalent according to the spec sheets, but perhaps more likely to die the day after the warranty expires. Substitution can be interpreted as “bait and switch”, so most vendors are scared of getting called on this. Very few will hold their position if you press the matter.

Another red flag: “Only warranted in supported environments”. This may mean they won't honor a warranty on a non-Windows system at all, or it may mean they'll insist on installing the Unix on disk themselves.

One absolute show-stopper is the phrase “All sales are final”. This means you have *no* options if a part doesn't work. Avoid any company with this policy.

Special Questions to Ask Web/Mail-Order Vendors Before Buying

- Does the vendor have the part or system presently in stock? Mail order companies tend to run with very lean inventories; if they don't have your item in stock, delivery may take longer. Possibly *much* longer.
- Does the vendor pay for shipping? What's the delivery wait?
- If you need to return your system, is there a restocking fee? and will the vendor cover the return freight? Knowing the restocking fee can be particularly important, as they make keep you from getting real satisfaction on a bad major part. Avoid dealing with anyone who quotes more than a 15% restocking fee — and it's a good idea, if possible, to avoid any dealer who charges a restocking fee at all.

Warranties are tricky. There are companies whose warranties are invalidated by opening the case. Some of those companies sell upgradeable systems, but only authorized service centers can do upgrades without invalidating the warranty. Sometimes a system is purchased with the warranty already invalidated. There are vendors who buy minimal systems and upgrade them with cheap RAM and/or disk drives. If the vendor is not an authorized service center, the manufacturer's warranty is invalidated. The only recourse in case of a problem is the vendor's warranty. So beware!

Payment Method

It's a good idea to pay with AmEx or Visa or MasterCard; that way you can stop payment if you get a lemon, and may benefit from a buyer-protection plan using the credit card company's clout (not all cards offer buyer-protection plans, and some that do have restrictions which may be applicable). However, watch for phrases like “Credit card surcharges apply” or “All prices reflect 3% cash discount” which mean you're going to get socked extra if you pay by card.

Note that many credit-card companies have clauses in their standard contracts forbidding such surcharges. You can (and should) report such practices to your credit-card issuer. If you already paid the surcharge, they will usually see to it that it is returned to you. Credit-card companies will often stop dealing with businesses that repeat such behavior.

Which Clone Vendors to Talk To

Some pans

Gateway: may also be a vendor to avoid. Apparently their newer machines don't have parity bits in their memories; memory is tested only on reboot. This is dubious design even for Windows, and totally unacceptable for Unix.

Some picks

In early August 2001 I designed an 'Ultimate Linux Box' with Gary Sandine and John Pearson of Los Alamos Computers [<http://lanm-pc.com>]; you can read all about it [<http://www.catb.org/~esr/writings/ultimate-linux-box/>] These guys know what they are doing and are fun to work with. If you need a high-end Linux workstation, or your laboratory needs a computer cluster, talk with them.

After You Take Delivery

Your configuration is custom and involves slightly unusual hardware. Therefore, keep a copy of the configuration you wrote down, and check it against the invoice and the actual delivered hardware. If there is a problem, calling back your vendor right away will maximize your chances of getting the matter settled quickly.

Software to go with your hardware

I used to maintain an entire separate FAQ on Unices for 386/486 and Pentium hardware. Times change, industries evolve, and I can now replace that FAQ with just three words:

Go get Linux!

Note

FreeBSD or OpenSolaris are currently niche choices, but if they offer something you need that Linux doesn't, don't let me stop you from trying one or both of them.

Other Resources on Building Linux PCs

The PC Tech Guide [<http://www.pctechguide.com/>] offers pretty comprehensive descriptions of PC hardware technologies.

The Caveat Emptor [<http://hawks.ha.md.us/hardware/>] guide has an especially good section on evaluating monitor specifications.

Anthony Olszewski's Assembling A PC [<http://www.computercraft.com/docs/pcbuild.html>] is an excellent guide to the perplexed. Not Linux-specific.

Tom's Hardware Guide [<http://www.sysdoc.pair.com/>] covers many hardware issues exhaustively. It is especially good about CPU chips and motherboards. Full of ads and slow-loading graphics, though.

The System Optimization Site [<http://www.sysopt.com>] has many links to other worthwhile sites for hardware buyers.

Christopher B. Browne has a page on Linux VARs [<http://linuxfinances.info/info/linuxvars.html#VARs>] that build systems. He also recommends the Linux VAR HOWTO [<http://en.tldp.org/HOWTO/VAR-HOWTO.html>].

There's a Building Your Own PC [<http://www.aplus.net/docs/facts/the-ultimate-guide-to-building-your-own-pc.htm>] page. It's more oriented towards building from parts than this one. Less technical depth in most areas, but better coverage of some including RAM, soundcards and motherboard installation. Features nifty and helpful graphics, one of the better graphics-intensive pages I've seen. However, the hardware-selection advice is out of date.

The Linux Hardware Database [<http://en.tldp.org/HOWTO/Hardware-HOWTO>] .

The Silent PC Reviews [<http://www.silentpcreview.com/>] site has lots of good material on building quiet PCs.