

# RAM

In this chapter, you will learn how to

- Identify the different types of RAM packaging
- Explain the varieties of DRAM
- Install RAM properly
- Perform basic RAM troubleshooting

Anytime someone comes up to me and starts professing their computer savvy, I ask them a few questions to see how much they really know. Just in case you and I ever meet and you decide you want to “talk tech” with me, I’ll tell you my first two questions just so you’ll be ready. Both involve *random access memory (RAM)*, the working memory for the CPU.

1. “How much RAM is in your computer?”
2. “What is RAM and why is it so important that every PC have some?”

Can you answer either of these questions? Don’t fret if you can’t—you’ll know how to answer both of them before you finish this chapter. Let’s start by reviewing what you know about RAM thus far.




**EXAM TIP** The CompTIA A+ certification domains use the term *memory* to describe the short-term storage used by the PC to load the operating system and running applications. The more common term in the industry is *RAM*, for random access memory, the kind of short-term memory you’ll find in every computer. More specifically, the primary system RAM is *dynamic random access memory (DRAM)*. For the most part, this book uses the terms RAM and DRAM.


When not in use, programs and data are held in mass storage, which usually means a hard drive, but could also mean a USB thumb drive, a CD-ROM, or some other device that can hold data when the computer is turned off (Figure 4-1). When you load a program by clicking an icon in Windows, the program is copied from the mass storage device to RAM and then run (Figure 4-2).

**Figure 4-1**

Mass storage holds unused programs.

**Figure 4-2**

Programs run in RAM.



You saw in the previous chapter that the CPU uses dynamic random access memory (DRAM) as RAM for all PCs. Just like CPUs, DRAM has gone through a number of evolutionary changes over the years, resulting in improved DRAM technologies with names such as SDRAM, RDRAM, and DDR RAM. This chapter starts by explaining how DRAM works, and then moves into the types of DRAM used over the last few years to see how they improve on the original DRAM. The third section, "Working with RAM," goes into the details of finding and installing RAM. The chapter finishes with troubleshooting RAM problems.

## Historical/Conceptual

### Understanding DRAM

As discussed in Chapter 3, DRAM functions like an electronic spreadsheet, with numbered rows containing cells and each cell holding a one or zero. Now let's look at what's physically happening. Each spreadsheet cell is a special type of semiconductor that can hold a single bit—one or zero—using microscopic capacitors and transistors. DRAM makers put these semiconductors into chips that can hold a certain number of bits. The bits inside the chips are organized in a rectangular fashion using rows and columns.

Each chip has a limit on the number of lines of code it can contain. Think of each line of code as one of the rows on the electronic spreadsheet; one chip might be able to store a million rows of code while another chip can store over a billion lines. Each chip also has a limit on the width of the lines of code it can handle, so one chip might

handle 8-bit-wide data while another might handle 16-bit-wide data. Techs describe chips by bits rather than bytes, so  $\times 8$  and  $\times 16$ , respectively. Just as you could describe a spreadsheet by the number of rows and columns—John’s accounting spreadsheet is huge, 48 rows  $\times$  12 columns—memory makers describe RAM chips the same way. An individual DRAM chip that holds 1,048,576 rows and 8 columns, for example, would be a  $1\text{ M} \times 8$  chip, with “M” as shorthand for “mega,” just like in megabytes ( $2^{20}$  bytes). It is difficult if not impossible to tell the size of a DRAM chip just by looking at it—only the DRAM makers know the meaning of the tiny number on the chips (although sometimes you can make a good guess). See Figure 4-3.

**Figure 4-3**

What do these numbers mean?




## Organizing DRAM

Due to its low cost, high speed, and capability to contain a lot of data in a relatively small package, DRAM has been the standard RAM used in all computers—not just PCs—since the mid-1970s. DRAM can be found in just about everything, from automobiles to automatic bread makers.

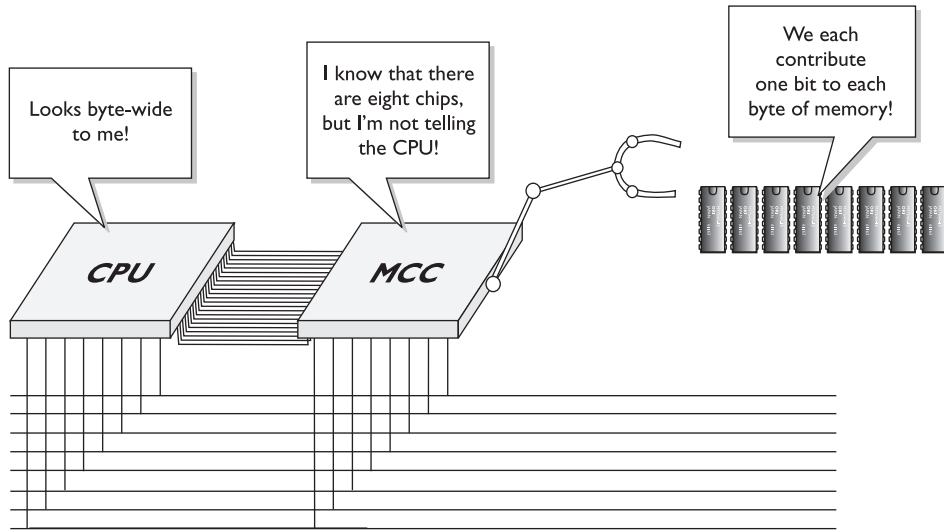
The PC has very specific requirements for DRAM. The original 8088 processor had an 8-bit frontside bus. All the commands given to an 8088 processor were in discrete, 8-bit chunks. Therefore, you needed RAM that could store data in 8-bit (1-byte) chunks, so that each time the CPU asked for a line of code, the memory controller could put an 8-bit chunk on the data bus. This optimized the flow of data into (and out from) the CPU.

Although today’s DRAM chips may have widths greater than 1 bit, back in the old days all DRAM chips were 1 bit wide. That means you only had sizes like  $64\text{ K} \times 1$  or  $256\text{ K} \times 1$ —always 1 bit wide. So how was 1-bit-wide DRAM turned into 8-bit-wide memory? The answer was quite simple: just take eight 1-bit-wide chips and electronically organize them with the memory controller chip to be eight wide. First, put eight 1-bit-wide chips in a row on the motherboard (Figure 4-4), and then wire up this row



**Figure 4-4** One row of DRAM

of DRAM chips to the memory controller chip (which has to be designed to handle this) to make byte-wide memory (Figure 4-5). You just made eight 1-bit-wide DRAM chips look like a single 8-bit-wide DRAM chip to the CPU.




**Figure 4-5** The MCC in action

## Practical DRAM

Okay, before you learn more about DRAM, I need to make a critical point extremely clear. When you first saw the 8088's machine language in the previous chapter, all the examples in the "codebook" were exactly 1-byte commands. Figure 4-6 shows the codebook again—see how all the commands are 1 byte?

Well, the reality is slightly different. Most of the 8088 machine language commands are 1 byte, but a few more complex commands need 2 bytes. For example, the command below tells the CPU to move 163 bytes "up the RAM spreadsheet" and run whatever command is there. Cool, eh?

```
1110100110100011
```




8088 External Data Bus Codebook	
LIGHTS	MEANING
10000000	The next line is a number, put it in the AX register
10010000	The next line is a number, put it in the BX register
10110000	Add AX to BX and put the result in AX
11000000	Place the value of AX on the External Data Bus
00000000	The number is 0
00000001	The number is 1
00000010	The number is 2
00000011	The number is 3
00000100	The number is 4
00000101	The number is 5
00000110	The number is 6
00000111	The number is 7
00001000	The number is 8
00001001	The number is 9

**Figure 4-6** Codebook again

The problem here is that the command is 2 bytes wide, not 1 byte wide! So how did the 8088 handle this? Simple—it just took the command 1 byte at a time. It took twice as long to handle the command because the MCC had to go to RAM twice, but it worked.

Okay, so if some of the commands are more than 1 byte wide, why didn't Intel make the 8088 with a 16-bit frontside bus? Wouldn't that have been better? Well, Intel did! Intel invented a CPU called the 8086. The 8086 actually predates the 8088 and was

absolutely identical to the 8088 except for one small detail—it had a 16-bit frontside bus. IBM could have used the 8086 instead of the 8088 and used 2-byte-wide RAM instead of 1-byte-wide RAM. Of course, they would have needed to invent a memory controller chip that handled that kind of RAM (Figure 4-7).




**Figure 4-7** Pumped up 8086 MCC at work

Why didn't Intel sell IBM the 8086 instead of the 8088? There were two reasons. First, nobody had invented an affordable MCC or RAM that handled two bytes at a time. Sure, chips were invented, but they were *expensive* and IBM didn't think that anyone would want to pay \$12,000 for a personal computer. So IBM bought the Intel 8088, not the Intel 8086, and all our RAM came in bytes. But as you might imagine, it didn't stay that way too long.

## DRAM Sticks


As CPU data bus sizes increased, so too did the need for RAM wide enough to fill the bus. The Intel 80386 CPU, for example, had a 32-bit data bus and thus the need for 32-bit-wide DRAM. Imagine having to line up 32 1-bit-wide DRAM chips on a motherboard. Talk about a waste of space! Figure 4-8 shows an example of motherboard RAM run amuck.



**Figure 4-8** That's a lot of real estate used by RAM chips!

DRAM manufacturers responded by creating wider DRAM chips, such as  $\times 4$ ,  $\times 8$ , and  $\times 16$ , and putting multiples of them on a small circuit board called a *stick* or *module*. Figure 4-9 shows an early stick, called a *single inline memory module (SIMM)*, with eight DRAM chips. To add RAM to a modern machine means you need to get the right stick or sticks for the particular motherboard. Your motherboard manual tells you precisely what sort of module you need and how much RAM you can install.

**Figure 4-9**  
A 72-pin SIMM



Modern CPUs are a lot smarter than the old Intel 8088. Their machine languages have some commands that are up to 64 bits (8 bytes) wide. They also have at least a 64-bit front-side bus that can handle more than just 8 bits. They don't want RAM to give them a puny 8 bits at a time! To optimize the flow of data into and out of the CPU, the modern MCC provides at least 64 bits of data every time the CPU requests information from RAM.



**NOTE** Some MCCs are 128 bits wide.

## Dealing with Old RAM

Often in the PC world, old technology and ways of doing things get re-implemented with some newer technology. Learning how things worked back in the ancient days can stand a tech in good stead. Perhaps more importantly, many thousands of companies—including hospitals, auto repair places, and more—use very old, proprietary applications that keep track of medical records, inventory, and so on. If you're called to work on one of these ancient systems, you need to know how to work with old parts, so here's a suggestion.

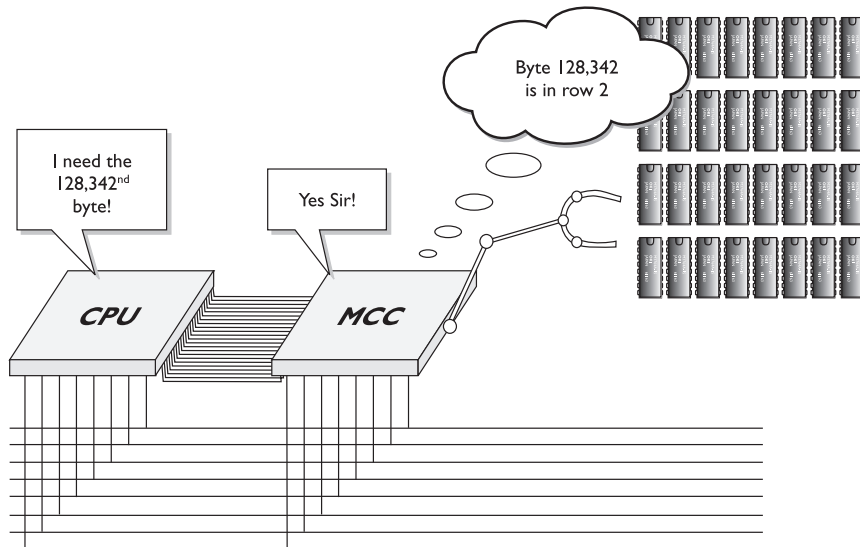
Obtain an old computer, such as an 80386 or 80486. Ask your uncle or cousin or Great Aunt Edna if they have a PC collecting dust in a closet that you can use. Failing that, go to a secondhand store or market and buy one for a few dollars.

Open up the system and check out the RAM. Remove the RAM from the motherboard and then replace it to familiarize yourself with the internals. You never know when some critical system will go down and need repair immediately—and you're the one to do it!

## Current RAM

Modern DRAM sticks come in 32-bit- and 64-bit-wide form factors with a varying number of chips. Many techs describe these memory modules by their width, so  $\times 32$  and  $\times 64$ . Note that this number does *not* describe the width of the individual DRAM chips on the module! When you read or hear about *by whatever* memory, simply note that you need to know whether that person is talking about the DRAM width or the module width.

When the CPU needs certain bytes of data, it requests those bytes via the address bus. The CPU does not know the physical location of the RAM that stores that data, nor the physical makeup of the RAM—such as how many DRAM chips work together to provide the 64-bit-wide memory rows. The MCC keeps track of this and just gives the CPU whichever bytes it requests (Figure 4-10).



**Figure 4-10** The MCC knows the real location of the DRAM.

## Consumer RAM

So if modern DRAM modules come in sizes much wider than a byte, why do people still use the word “byte” to describe how much DRAM they have? Convention. Habit. Rather than using a label that describes the electronic structure of RAM, common usage describes the *total capacity of RAM on a stick in bytes*. John has a single 512-MB stick of RAM on his motherboard, for example, while Sally has two 256-MB sticks. Both systems have a total of 512 MB of system RAM. That’s what your clients care about, after all, because having enough RAM makes their systems snappy and stable; not enough and the systems run poorly. As a tech, you need to know more, of course, to pick the right RAM for many different types of computers.

## Essentials

### Types of RAM

Development of newer, wider, and faster CPUs and MCCs motivate DRAM manufacturers to invent new DRAM technologies that deliver enough data at a single pop to optimize the flow of data into and out of the CPU.



**EXAM TIP** Old RAM—really old RAM—was called *fast page mode (FPM)* RAM. This ancient RAM used a totally different technology that was not tied to the system clock. If you ever hear of FPM RAM, it’s going to be in a system that’s over a decade old. Be careful! CompTIA likes to use older terms like this to throw you off!

### SDRAM

Most modern systems use some form of *synchronous DRAM (SDRAM)*. SDRAM is still DRAM, but it is *synchronous*—tied to the system clock, just like the CPU and MCC, so the MCC knows when data is ready to be grabbed from SDRAM. This results in little wasted time.


SDRAM made its debut in 1996 on a stick called a *dual inline memory module (DIMM)*. The early SDRAM DIMMs came in a wide variety of pin sizes. The most common pin sizes found on desktops were the 168-pin variety. Laptop DIMMs came in 68-pin, 144-pin

**Figure 4-11**  
144-pin micro-DIMM  
(photo courtesy of  
Micron Technology,  
Inc.)



(Figure 4-11), or 172-pin *micro-DIMM* packages; and the 72-pin, 144-pin, or 200-pin *small outline DIMM (SO DIMM)* form factors (Figure 4-12). With the exception of the 32-bit 72-pin SO DIMM, all these DIMM varieties delivered 64-bit-wide data to match the 64-bit data bus of every CPU since the Pentium.

**Figure 4-12**  
A 168-pin DIMM  
above a 144-pin  
SO-DIMM



To take advantage of SDRAM, you needed a PC designed to use SDRAM. If you had a system with slots for 168-pin DIMMs, for example, your system used SDRAM. A DIMM in any one of the DIMM slots could fill the 64-bit bus, so each slot was called a *bank*. You could install one, two, or more sticks and the system would work. Note that on laptops that used the 72-pin SO DIMM, you needed to install two sticks of RAM to make a full bank because each stick only provided half the bus width.

SDRAM tied to the system clock, so it had a clock speed that matched the frontside bus. Five clock speeds were commonly used on the early SDRAM systems: 66, 75, 83, 100, and 133 MHz. The RAM speed had to match or exceed the system speed or the computer would be unstable or wouldn't work at all. These speeds were prefixed with a "PC" in the front based on a standard forwarded by Intel, so SDRAM speeds were PC66 through PC133. For a Pentium III computer with a 100-MHz frontside bus, you needed to buy SDRAM DIMMs rated to handle it, such as PC100 or PC133.

## RDRAM

When Intel was developing the Pentium 4, they knew that regular SDRAM just wasn't going to be fast enough to handle the quad-pumped 400-MHz frontside bus. Intel announced plans to replace SDRAM with a very fast, new type of RAM developed by Rambus, Inc. called Rambus DRAM, or simply *RDRAM* (Figure 4-13). Hailed by Intel as the next great leap in DRAM technology, RDRAM could handle speeds up to 800 MHz, which gave Intel plenty of room to improve the Pentium 4.

**Figure 4-13**  
RDRAM




**NOTE** The 400-MHz frontside bus speed wasn't achieved by making the system clock faster—it was done by making CPUs and MCCs capable of sending 64 bits of data two or four times for every clock cycle, effectively doubling or quadrupling the system bus speed.

RDRAM was greatly anticipated by the industry for years, but industry support for RDRAM proved less than enthusiastic due to significant delays in development, plus a price many times that of SDRAM. Despite this grudging support, almost all major PC makers sold systems that used RDRAM—for a while. From a tech's standpoint, RDRAM shares almost all of the characteristics of SDRAM. A stick of RDRAM is called a *RIMM*. In this case, however, the letters don't actually stand for anything; they just rhyme: SIMMs, DIMMs, and now RIMMs, get it?

RDRAM RIMMs came in two sizes: a 184-pin for desktops and a 160-pin SO-RIMM for laptops. RIMMs were keyed differently from DIMMs to ensure that even though they are the same basic size, you couldn't accidentally install a RIMM in a DIMM slot or vice versa. RDRAM also had a speed rating: 600 MHz, 700 MHz, 800 MHz, or 1066 MHz. RDRAM employed an interesting dual-channel architecture. Each RIMM was 64 bits wide, but the Rambus MCC alternated between two sticks to increase the speed of data retrieval. You were required to install RIMMs in pairs to use this dual-channel architecture.

RDRAM motherboards also required that all RIMM slots be populated. Unused pairs of slots needed a passive device called a *continuity RIMM* (CRIMM) installed in each slot to enable the RDRAM system to terminate properly. Figure 4-14 shows a CRIMM.

**Figure 4-14**  
CRIMM



RDRAM offered dramatic possibilities for high-speed PCs, but ran into three roadblocks that betamaxed it. First, the technology was owned wholly by Rambus—if you wanted to make it, you had to pay the licensing fees they charged. That led directly to the second problem, expense. RDRAM cost substantially more than SDRAM. Third, Rambus and Intel made a completely closed deal for the technology. RDRAM worked only on Pentium 4 systems using Intel-made MCCs. AMD was out of luck. Clearly, the rest of the industry had to look for another high-speed RAM solution.



**NOTE** *Betamaxed* is slang for “made it obsolete because no one bought it, even though it was a superior technology to the winner in the marketplace.” Refers to the VHS versus Betamax wars in the old days of video cassette recorders.

## DDR SDRAM


AMD and many major system and memory makers threw their support behind *double data rate SDRAM* (*DDR SDRAM*). DDR SDRAM basically copied Rambus, doubling the throughput of SDRAM by making two processes for every clock cycle. This synchronized (pardon the pun) nicely with the Athlon and later AMD processors' double-

pumped frontside bus. DDR SDRAM could not run as fast as RDRAM—although relatively low frontside bus speeds made that a moot point—but cost only slightly more than regular SDRAM.




**NOTE** Most techs drop some or all of the SDRAM part of DDR SDRAM when engaged in normal geekspeak. You'll hear the memory referred to as DDR, DDR RAM, and the weird hybrid, DDRAM.

DDR SDRAM for desktops comes in 184-pin DIMMs. These DIMMs match 168-pin DIMMs in physical size, but not in pin compatibility (Figure 4-15). The slots for the two types of RAM appear similar as well, but have different guide notches, making it impossible to insert either type of RAM into the other's slot. DDR SDRAM for laptops comes in either 200-pin SO-DIMMs or 172-pin micro-DIMMs (Figure 4-16).



**Figure 4-15** DDR SDRAM

**Figure 4-16**  
172-pin DDR  
SDRAM micro-  
DIMM (photo  
courtesy of  
Kingston/Joint  
Harvest)



**NOTE** RAM makers use the term *single data rate SDRAM (SDR SDRAM)* for the original SDRAM to differentiate it from DDR SDRAM.

DDR sticks use a rather interesting naming convention—actually started by the Rambus folks—based on the number of bytes per second of data throughput the RAM can handle. To determine the bytes per second, take the MHz speed and multiply by 8 bytes (the width of all DDR SDRAM sticks). So 400 MHz multiplied by 8 is 3200 bytes

per second. Put the abbreviation “PC” in the front to make the new term: PC3200. Many techs also use the naming convention used for the individual DDR chips; for example, *DDR400* refers to a 400-MHz DDR SDRAM chip running on a 200-MHz clock. Even though the term *DDRxxx* is really just for individual DDR chips and the term *PCxxxx* is for DDR sticks, this tradition of two names for every speed of RAM is a bit of a challenge as both terms are commonly used interchangeably. Table 4-1 shows all the speeds for DDR—not all of these are commonly used.


<b>Clock Speed</b>	<b>DDR Speed Rating</b>	<b>PC Speed Rating</b>
100 MHz	DDR200	PC1600
133 MHz	DDR266	PC2100
166 MHz	DDR333	PC2700
200 MHz	DDR400	PC3200
217 MHz	DDR433	PC3500
233 MHz	DDR466	PC3700
250 MHz	DDR500	PC4000
275 MHz	DDR550	PC4400
300 MHz	DDR600	PC4800

**Table 4-1** DDR Speeds

Following the lead of AMD, VIA, and other manufacturers, the PC industry adopted DDR SDRAM as the standard system RAM. Intel relented and stopped producing motherboards and memory controllers that required RDRAM in the summer of 2003.

There’s one sure thing about PC technologies—any good idea that can be copied will be copied. One of Rambus’ best concepts was the dual-channel architecture—using two sticks of RDRAM together to increase throughput. Manufacturers have released motherboards with MCCs that support dual-channel architecture using DDR SDRAM. Dual-channel DDR motherboards use regular DDR sticks, although manufacturers often sell RAM in matched pairs, branding them as dual-channel RAM.

Dual-channel DDR works like RDRAM in that you must have two identical sticks of DDR and they must snap into two paired slots. Unlike RDRAM, dual-channel DDR doesn’t have anything like CRIMMS—you don’t need to put anything into unused slot pairs. Dual-channel DDR technology is very flexible, but also has a few quirks that vary with each system. Some motherboards have three DDR SDRAM slots, but the dual-channel DDR works only if you install DDR SDRAM in two of the slots (Figure 4-17). If you populate the third slot, the system will use the full capacity of RAM installed, but turns off the dual-channel feature—and no, it doesn’t tell you! The dual slots are blue; the third slot is black, which you could clearly see if this weren’t a black-and-white photo.




**Figure 4-17** An nForce motherboard showing the three RAM slots. The two slots bracketing the slim space can run as dual channel as long as you don't populate the third slot.

## DDR2

The fastest versions of DDR RAM run at a blistering PC4800. That's 4.8 gigabytes per second (GBps) of data throughput! You'd think that kind of speed would satisfy most users, and to be honest, DRAM running at approximately 5 GBps really is plenty fast—for now. However, the ongoing speed increases ensure that even these speeds won't be good enough in the future. Knowing this, the RAM industry came out with DDR2, the successor to DDR. *DDR2* is DDR RAM with some improvements in its electrical characteristics, enabling it to run even faster than DDR while using less power. The big speed increase from DDR2 comes by clock doubling the input/output circuits on the chips. This does not speed up the core RAM—the part that holds the data—but speeding up the input/output and adding special buffers (sort of like a cache) makes DDR2 runs much faster than regular DDR. DDR2 uses a 240-pin DIMM that's not compatible with DDR. You'll find motherboards running both single-channel and dual-channel DDR2 (Figure 4-18).

**Figure 4-18**  
240-pin DDR2  
DIMM



The following table shows some of the common DDR2 speeds.

Core RAM Clock Speed	DDR I/O Speed	DDR2 Speed Rating	PC Speed Rating
100 MHz	200 MHz	DDR2-400	PC2-3200
133 MHz	266 MHz	DDR2-533	PC2-4200
166 MHz	333 MHz	DDR2-667	PC2-5300
200 MHz	400 MHz	DDR2-800	PC2-6400
250 MHz	500 MHz	DDR2-1000	PC2-8000


## RAM Variations

Within each class of RAM, you'll find variations in packaging, speed, quality, and the capability to handle data with more or fewer errors. Higher-end systems often need higher-end RAM, so knowing these variations is of crucial importance to techs.

### Double-Sided DIMMs

Every type of RAM stick, starting with the old FPM SIMMs and continuing through to 240-pin DDR2 SDRAM, comes in one of two types: single-sided and double-sided. As their name implies, single-sided sticks only have chips on one side of the stick. Double-sided sticks have chips on both sides (Figure 4-19). The vast majority of RAM sticks are single-sided, but there are plenty of double-sided sticks out there. Double-sided sticks are basically two sticks of RAM soldered onto one board. There's nothing wrong with double-sided RAM other than the fact that some motherboards either can't use them or can only use them in certain ways—for example, only if you use a single stick and it goes into a certain slot.

**Figure 4-19**  
Double-sided DDR  
SDRAM




## Latency

If you've shopped for RAM lately, you may have noticed terms such as "CL2" or "low-latency" as you tried to determine which RAM to purchase. You might find two otherwise identical RAM sticks with a 20 percent price difference and a salesperson pressuring you to buy the more expensive one because it's "faster" even though both sticks say DDR400 (Figure 4-20).

**Figure 4-20**

Why is one more expensive than the other?



RAM responds to electrical signals at varying rates. When the memory controller starts to grab a line of memory, for example, there's a slight delay; think of it as the RAM getting off the couch. After the RAM sends out the requested line of memory, there's another slight delay before the memory controller can ask for another line—the RAM sat back down. The delay in RAM's response time is called its *latency*. RAM with a lower latency—such as CL2—is faster than RAM with a higher latency—such as CL3—because it responds more quickly.

Latency numbers reflect how many clicks of the system clock it takes before the RAM responds. If you speed up the system clock, say from 166 MHz to 200 MHz, the same stick of RAM might take an extra click before it can respond. When you take RAM out of an older system and put it into a newer one, you might get a seemingly dead PC, even though the RAM fits in the DIMM slot. Many motherboards enable you to adjust the RAM timings manually. If so, try raising the latency to give the slower RAM time to respond. See Chapter 5, "BIOS and CMOS," to learn how to make these adjustments (and how to recover if you make a mistake!).

From a tech's standpoint, you need to get the proper RAM for the system on which you're working. If you put a high latency stick in a motherboard set up for a low latency stick, you'll get an unstable or completely dead PC. Check the motherboard manual and get the quickest RAM the motherboard can handle and you should be fine.




**NOTE** CAS stands for *column array strobe*, one of the wires (along with the *row array strobe*) in the RAM that helps the memory controller find a particular bit of memory. Each of these wires require electricity to charge up before they can do their job. This is one of the aspects of latency.

## Parity and ECC

Given the high speeds and phenomenal amount of data moved by the typical DRAM chip, it's possible that a RAM chip might occasionally give bad data to the memory controller. This doesn't necessarily mean that the RAM has gone bad. This could be an occasional hiccup caused by some unknown event that makes a good DRAM chip say a bit is a zero when it's really a one. In most cases, you won't even notice when such a rare event happens. In some environments, however, even these rare events are intolerable. A bank server handling thousands of online transactions per second, for example, can't risk even the smallest error. These important computers need a more robust, fault-resistant RAM.

The first type of error-detecting RAM was known as parity RAM (Figure 4-21). *Parity* RAM stored an extra bit of data (called the parity bit) that the MCC used to verify if the data was correct. Parity wasn't perfect—it wouldn't always detect an error, and if the MCC did find an error, it couldn't correct the error. For years, parity was the only available way to tell if the RAM made a mistake.

**Figure 4-21**  
Ancient parity  
RAM stick



Today's PCs that need to watch for RAM errors use a special type of RAM called *error correction code (ECC)* RAM. ECC is a major advance in error checking on DRAM. First, ECC detects any time a single bit is incorrect. Second, ECC fixes these errors on the fly. The checking and fixing come at a price, as ECC RAM is always slower than non-ECC RAM.




**NOTE** Some memory manufacturers call the technology *Error Checking and Correction (ECC)*. Don't be thrown off if you see the phrase—it's the same thing, just a different marketing slant for error correction code.

ECC DRAM comes in every DIMM package type. You might be tempted to say "Gee, maybe I want to try this ECC RAM!" Well, don't! To take advantage of ECC RAM, you need a motherboard with an MCC designed to use ECC. Only expensive motherboards for high-end systems use ECC. The special-use-only nature of ECC makes it fairly rare. There are plenty of techs out there with years of experience who've never even seen ECC RAM.

## Buffered/Registered DRAM

Your average PC motherboard accepts no more than four sticks of DRAM because more than four physical slots for sticks gives motherboard designers some serious electrical headaches. Yet some systems that use a lot of RAM need the capability to use more DRAM sticks on the motherboard, often six or eight. To get around the electrical hassles, special DRAM sticks add a buffering chip to the stick that acts as an intermediary between the DRAM and the MCC. These special DRAMs are called *buffered* or *registered* DRAM (Figure 4-22). The DDR2 version is called *fully buffered*. (See the “Beyond A+” section for more details.)

**Figure 4-22**  
Buffered RAM




Like ECC, you must have a motherboard with an MCC designed to use this type of DRAM. Rest assured that such a motherboard will have a large number of RAM slots. Buffered/registered RAM is rare (maybe not quite as rare as ECC RAM), and you’ll never see it in the typical desktop system.

## Working with RAM

Whenever someone comes up to me and asks what single hardware upgrade they can do to improve their system performance, I always tell them the same thing—add more RAM. Adding more RAM can improve overall system performance, processing speed, and stability—if you get it right. Botching the job can cause dramatic system instability, such as frequent, random crashes and reboots. Every tech needs to know how to install and upgrade system RAM of all types.

To get the desired results from a RAM upgrade, you must first determine if insufficient RAM is the cause of system problems. Second, you need to pick the proper RAM for the system. Finally, you must use good installation practices. Always store RAM sticks in anti-static packaging whenever they’re not in use, and use strict ESD handling procedures. Like many other pieces of the PC, RAM is *very* sensitive to ESD and other technician abuse (Figure 4-23)!

**Figure 4-23**  
Don't do this!  
Grabbing the  
contacts is a  
*bad thing.*



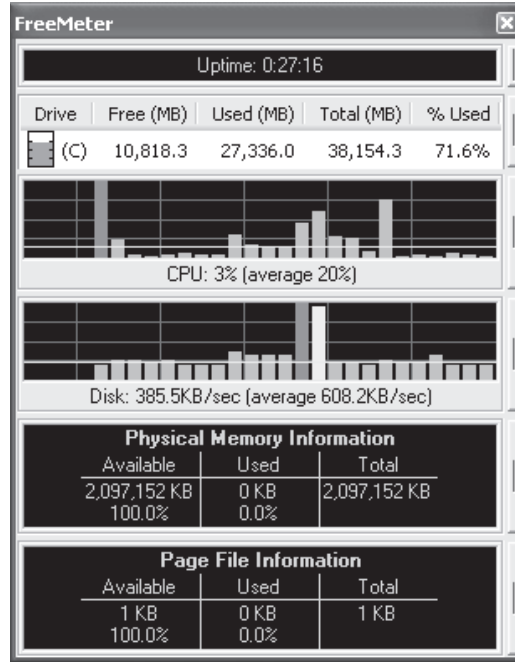
## Do You Need RAM?

Two symptoms point to the need for more RAM in a PC: general system sluggishness and excessive hard drive accessing. If programs take forever to load and running programs seem to stall and move more slowly than you would like, the problem could stem from insufficient RAM. A friend with a new Windows XP Professional system complained that her PC seemed snappy when she first got it, but takes a long time to do the things she wants to do with it, such as photograph retouching in Adobe Photoshop and document layout for a print zine she produces. Her system had only 256 MB of RAM, sufficient to run Windows XP Professional, but woefully insufficient for her tasks—she kept maxing out the RAM and thus the system slowed to a crawl. I replaced her stick with a pair of 1-GB sticks and suddenly she had the powerhouse workstation she desired.

Excessive hard drive activity when you move between programs points to a need for more RAM. Every Windows PC has the capability to make a portion of your hard drive look like RAM in case you run out of real RAM. This is called the *page file* or *swap file*. This book discusses page files in Chapter 15, but for now appreciate the fact that if you fill your RAM up with programs, your PC will automatically start loading some programs into the page file. You can't see this process taking place just by looking at the screen—these swaps are done in the background. But you will notice the hard drive access LED going crazy as Windows rushes to move programs between RAM and the page file in a process called *disk thrashing*. Windows uses the page file all the time, but excessive disk thrashing suggests that you need more RAM.

You can diagnose excessive disk thrashing through simply observing the hard drive access LED flashing, or through various third-party tools. I like FreeMeter ([www.tiler.com/freemeter/](http://www.tiler.com/freemeter/)). It's been around for quite a while, runs on all versions of Windows, and is easy to use (Figure 4-24). Notice on the FreeMeter screenshot that some amount of the page file is being used. That's perfectly normal.

**Figure 4-24**  
FreeMeter



### System RAM Recommendations

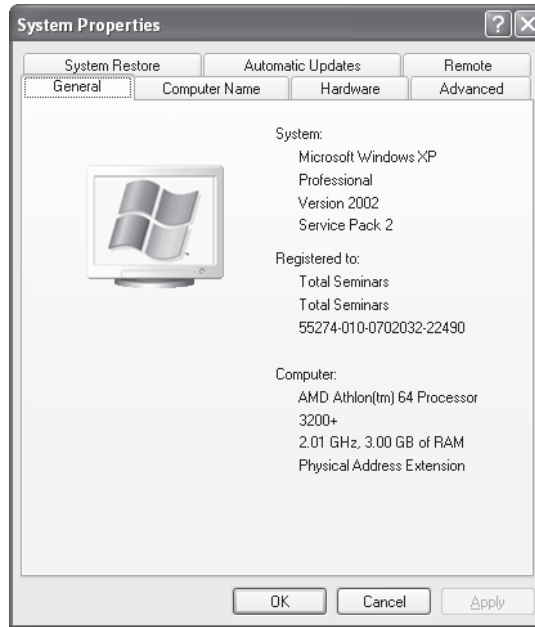
Microsoft sets the minimum RAM requirements listed for the various Windows operating systems very low to get the maximum number of users to upgrade or convert, and that's fine. A Windows XP Professional machine will run well enough on 128 MB of RAM. Just don't ask it to do any serious computing, like run Doom III! Here are my recommendations for system RAM.

Operating System	Reasonable Minimum	Solid Performance	Power User
Windows 2000	128 MB	256 MB	512 MB
Windows XP	256 MB	512 MB	1 GB
Windows Vista	512 MB	1 GB	2 GB

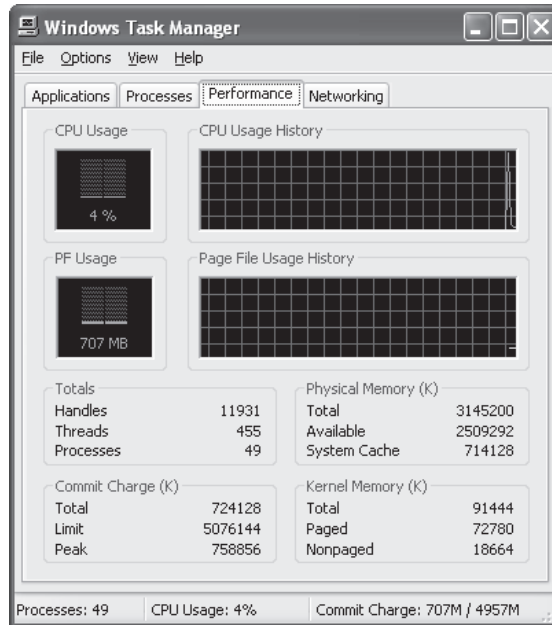
### Determining Current RAM Capacity

Before you go get RAM, you obviously need to know how much RAM you currently have in your PC. Every version of Windows works the same way. Just select the Properties for My Computer to see how much RAM is in your system (Figure 4-25). If you have a newer keyboard, you can access the screen with the WINDOWS-PAUSE/BREAK keystroke combination. Windows 2000 and XP come with the handy Performance tab under the Task Manager (as shown in Figure 4-26).

**Figure 4-25**  
Mike has a lot  
of RAM!



**Figure 4-26**  
Performance tab on  
the Windows XP  
Task Manager



## Getting the Right RAM

To do the perfect RAM upgrade, determine the optimum capacity of RAM to install and then get the right RAM for the motherboard. Your first two steps toward these goals are the inside of the case and your motherboard manual. Open the case to see how many

sticks of RAM you have installed currently and how many free slots you have open. Check the motherboard book to determine the total capacity of RAM the system can handle and what specific technology works with your system. You can't put DDR SDRAM into a system that can only handle SDR SDRAM, after all, and it won't do you much good to install a pair of 512-MB DIMMs when your system tops out at 784 MB! Figure 4-27 shows the RAM limits for my Gigabyte motherboard.


**Dual Channel Memory Configuration**  
 The GA-K8N Ultra-SLI/GA-K8N Pro-SLI supports the Dual Channel Technology. When the Dual Channel Technology is activated, the bandwidth of memory bus will be double the original one.


Due to CPU limitation, if you want to operate the Dual Channel Technology, please follow the guidelines below for Dual Channel memory configuration.

1. Dual Channel mode will not be enabled if only one DDR memory module is installed.
2. To enable Dual Channel mode with 2 memory modules (it is recommended to use memory modules of identical brand, size, chips, and speed), you must install them into DIMM sockets of the same color.
3. To enable Dual Channel mode with 4 memory modules, it is recommended to use memory modules of identical brand, size, chips, and speed.

The following is a Dual Channel Memory configuration table: (DS: Double Side, SS: Single Side)

	DDR1	DDR2	DDR3	DDR4
2 memory modules	DS/SS	DS/SS	X	X
	X	X	DS/SS	DS/SS
4 memory modules	DS/SS	DS/SS	DS/SS	DS/SS

 **NOTE** If two memory modules are to be used to achieve Dual Channel mode, we recommend installing them in DDR1 and DDR2 DIMM sockets.

 **CAUTION** All of the memory configurations below will cause system unable to boot. (DS: Double Side, SS: Single Side)

	DDR1	DDR2	DDR3	DDR4
1 memory module	X	DS/SS	X	X
	X	X	X	DS/SS
2 memory modules	X	DS/SS	DS/SS	X
	DS/SS	X	X	DS/SS
	X	DS/SS	X	DS/SS
3 memory modules	DS/SS	DS/SS	DS/SS	X
	X	DS/SS	DS/SS	DS/SS
	DS/SS	X	DS/SS	DS/SS
	DS/SS	DS/SS	X	DS/SS

**Figure 4-27** The motherboard book shows how much RAM that Athlon 64 can handle.

### Mix and Match at Your Peril

All motherboards can handle different capacities of RAM. If you have three slots, you may put a 512-MB stick in one and a 1-GB stick in the other with a high chance of success. To ensure maximum stability in a system, however, shoot for as close as you can get to uniformity of RAM. Choose RAM sticks that match in technology, capacity, and speed. Even on motherboards that offer slots for radically different RAM types, I recommend uniformity.

### Mixing Speeds

With so many different DRAM speeds available, you may often find yourself tempted to mix speeds of DRAM in the same system. Although in many situations you can get away with mixing speeds on a system, the safest, easiest rule to follow is to use the speed of


DRAM specified in the motherboard book, and make sure that every piece of DRAM runs at that speed. In a worst-case scenario, mixing DRAM speeds can cause the system to lock up every few seconds or every few minutes. You might also get some data corruption. Mixing speeds sometimes works fine, but don't do your income tax on a machine with mixed DRAM speeds until the system has proven to be stable for a few days. The important thing to note here is that you won't break anything, other than possibly data, by experimenting.

Okay, I have mentioned enough disclaimers. Modern motherboards provide some flexibility regarding RAM speeds and mixing. First, you can use RAM that is faster than what the motherboard specifies. For example, if the system needs PC3200 DDR2 SDRAM, you may put in PC4200 DDR2 SDRAM and it should work fine. Faster DRAM is not going to make the system run any faster, however, so don't look for any system improvement.

Second, you can sometimes get away with putting one speed of DRAM in one bank and another speed in another bank, as long as all the speeds are as fast or faster than the speed specified by the motherboard. Don't bother trying to put different-speed DRAMs in the same bank with a motherboard that uses dual-channel DDR. Yes, it works once in a while, but it's too chancy. I avoid it.

## Installing DIMMs and RIMMs

Installing DRAM is so easy that it's one of the very few jobs I recommend to non-techie folks. First, attach an anti-static wrist strap or touch some bare metal on the power supply to ground yourself and avoid ESD. Then swing the side tabs on the RAM slots down from the upright position. Pick up a stick of RAM—don't touch those contacts! A good hard push down is usually all you need to ensure a solid connection. Make sure that the DIMM snaps into position to show it is completely seated. You will also notice that the two side tabs will move in to reflect a tight connection (Figure 4-28).




**Figure 4-28** Inserting a DIMM

## SPD

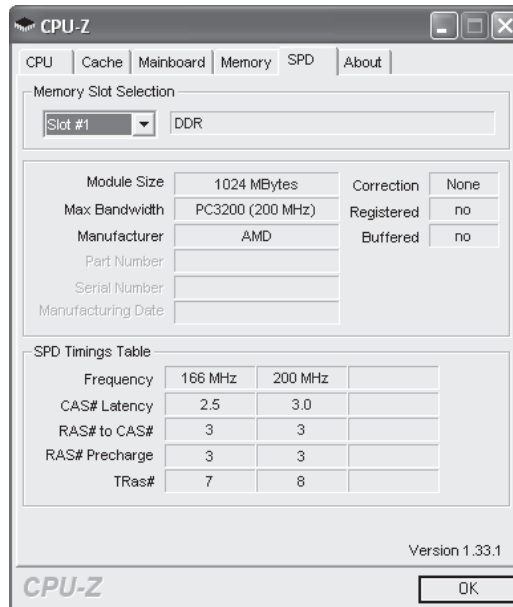
Your motherboard should detect and automatically set up any DIMM or RIMM you install, assuming you have the right RAM for the system. RAM makers add a handy chip to modern sticks called the *serial presence detect (SPD)* chip (Figure 4-29). The SPD stores all the information about your DRAM, including size, speed, ECC or non-ECC, registered or unregistered, and a number of other more technical bits of information.

When a PC boots, it queries the SPD so that the MCC knows how much RAM is on the stick, how fast it runs, and other information. Any program can query the SPD. Take a look at Figure 4-30 with the results of the popular CPU-Z program showing RAM information from the SPD.

**Figure 4-29**  
SPD chip on stick



**Figure 4-30**  
CPU-Z showing  
RAM information



All new systems count on the SPD to set the RAM timings properly for your system when it boots. If you add a RAM stick with a bad SPD, you'll get a POST error message and the system will not boot. You can't fix a broken SPD; you just buy a new stick of RAM.

## The RAM Count

After installing the new RAM, turn on the PC and watch the boot process closely. If you installed the RAM correctly, the RAM count on the PC will reflect the new value. If the RAM value stays the same, you probably have installed the RAM in a slot the motherboard doesn't want you to use (for example, if you need to use a particular slot first) or the RAM is not properly installed. If the computer does not boot and you've got a blank screen, you probably have not installed all the RAM sticks correctly. Usually, a good second look is all you need to determine the problem (Figures 4-31 and 4-32). Reseat or reinstall the RAM stick and try again.


**Figure 4-31**  
Hey, where's the rest of my RAM?!

```
Award Modular BIOS v6.00PG, An Energy Star Ally
Copyright (C) 1984-2005, Award Software, Inc.

GA-K8NP F13

Processor : AMD Athlon(tm) 64 Processor 3200+
<CPUID:0000F4A Patch ID:003A>
Memory Testing : 1048576K OK
CPU clock frequency : 200 Mhz

Detecting IDE drives ...
```



**Figure 4-32**  
RAM count after proper insertion of DIMMs

```
Award Modular BIOS v6.00PG, An Energy Star Ally
Copyright (C) 1984-2005, Award Software, Inc.

GA-K8NP F13

Processor : AMD Athlon(tm) 64 Processor 3200+
<CPUID:0000F4A Patch ID:003A>
Memory Testing : 3145728K OK
CPU clock frequency : 200 Mhz

Detecting IDE drives ...
```

RAM counts are confusing because RAM uses megabytes and gigabytes as opposed to millions and billions. Here are some examples of how different systems would show 256 MB of RAM:

268435456 (exactly 256 x 1 MB)

256M (some PCs try to make it easy for you)

262,144 (number of KB)


You should know how much RAM you're trying to install and use some common sense. If you've got 512 MB and you add another 512-MB stick, you need a number that looks like one gigabyte. If after you add the second stick, you see a RAM count of 524582912—that sure looks like 512 MB, not the one gigabyte!

## Installing SO-DIMMs in Laptops


It wasn't that long ago that adding RAM to a laptop was either impossible or required you to send the system back to the manufacturer. For years, every laptop maker had custom-made, proprietary RAM packages that were difficult to handle and staggeringly expensive. The wide acceptance of SO-DIMMs over the last few years has virtually erased these problems. All laptops now provide relatively convenient access to their SO-DIMMs, enabling easy replacement or addition of RAM.

Access to RAM usually requires removing a panel or lifting up the keyboard—the procedure varies among laptop manufacturers. Figure 4-33 shows a typical laptop RAM access panel. You can slide the panel off to reveal the SO-DIMMs. SO-DIMMs usually insert exactly like the old SIMMs—slide the pins into position and snap the SO-DIMM down into the retaining clips (Figure 4-34).

**Figure 4-33**  
A RAM access panel  
on a laptop



**Figure 4-34**  
Snapping in an  
SO-DIMM



Before doing any work on a laptop, turn the system off, disconnect it from the AC wall socket, and remove all batteries. Use an anti-static wrist strap because laptops are far more susceptible to ESD than desktop PCs.

## IT Technician

### Troubleshooting RAM

“Memory” errors show up in a variety of ways on modern systems, including parity errors, ECC error messages, system lockups, page faults, and other error screens in Windows. These errors can indicate bad RAM, but often point to something completely unrelated to RAM.

This is especially true with intermittent problems. The challenge for techs is to recognize these errors and then determine which part of the system caused the memory error.

You can get two radically different types of parity errors: real and phantom. Real parity errors are simply errors that the MCC detects from the parity or ECC chips (if you have them). The operating system then reports the problem in an error message, such as "Parity error at `xxxx:xxxxxxxx`," where `xxxx:xxxxxxxx` is a hexadecimal value (a string of numbers and letters, like `A5F2:004EEAB9`). If you get an error like this, write down the value. A real parity/ECC error will show up at the same place in memory each time, and almost always indicates that you have a bad RAM stick.

Phantom parity errors show up on systems that don't have parity or ECC memory. If Windows generates parity errors with different addresses, you most likely do *not* have a problem with RAM. These phantom errors can occur for a variety of reasons, including software problems, heat or dust, solar flares, fluctuations in the Force ... you get the idea.

System lockups and page faults (they often go hand in hand) in Windows can indicate a problem with RAM. A system lockup is when the computer stops functioning. A *page fault* is a milder error that can be caused by memory issues, but not necessarily system RAM problems. Certainly page faults *look* like RAM issues because Windows generates frightening error messages filled with long strings of hexadecimal digits, such as "KRNL386 caused a page fault at `03F2:25A003BC`." (See Figure 4-35.) Just because the error message contains a memory address, however, does not mean that you have a problem with your RAM. Write down the address. If it repeats in later error messages, you probably have a bad RAM stick. If Windows displays different memory locations, you need to look elsewhere for the culprit.

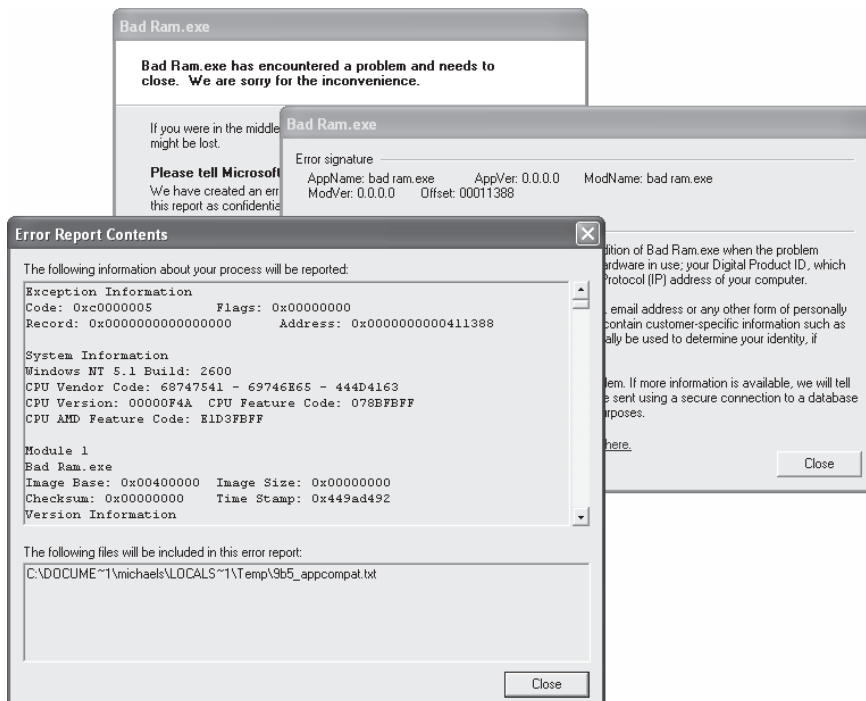


Figure 4-35 Windows error message

Every once in a while, something potentially catastrophic happens within the PC, some little electron hits the big red panic button, and the operating system has to shut down certain functions running in order to save data. This panic button inside the PC is called a *non-maskable interrupt (NMI)*, more simply defined as an interruption the CPU cannot ignore. An NMI manifests to the user as what techs lovingly call the *Blue Screen of Death (BSoD)*—a bright blue screen with a scary-sounding error message on it (Figure 4-36).

```
A problem has been detected and windows has been shut down to prevent damage
to your computer.

The problem seems to be caused by the following file: SPCMDCON.SYS

PAGE_FAULT_IN_NONPAGED_AREA

If this is the first time you've seen this Stop error screen,
restart your computer. If this screen appears again, follow
these steps:

Check to make sure any new hardware or software is properly installed.
If this is a new installation, ask your hardware or software manufacturer
for any windows updates you might need.

If problems continue, disable or remove any newly installed hardware
or software. Disable BIOS memory options such as caching or shadowing.
If you need to use Safe Mode to remove or disable components, restart
your computer, press F8 to select Advanced Startup Options, and then
select Safe Mode.

Technical information:

*** STOP: 0x00000050 (0xFD3094C2, 0x00000001, 0xFBFE7617, 0x00000000)

*** SPCMDCON.SYS - Address FBFE7617 base at FBFE5000, DateStamp 3d6dd67c
```

**Figure 4-36** Blue Screen of Death

Bad RAM sometimes triggers an NMI, although often the culprit lies with buggy programming or clashing code. The BSoD varies according to the operating system, and it would require a much lengthier tome than this one to cover all the variations. Suffice it to say that RAM *could* be the problem when that delightful blue screen appears.

Finally, intermittent memory errors can come from a variety of sources, including a dying power supply, electrical interference, buggy applications, buggy hardware, and so on. These errors show up as lockups, general protection faults, page faults, and parity errors, but never have the same address or happen with the same applications. Try the power supply first with non-application-specific intermittent errors of any sort.



**NOTE** A *general protection fault (GPF)* is an error that can cause an application to crash. Often, they're caused by programs stepping on each others' toes. Chapter 15 goes into more detail on GPFs and other Windows errors.

## Testing RAM

Once you discover that you may have a RAM problem, you have a couple of options. First, several companies manufacture hardware RAM testing devices, but unless you have a lot of disposable income, they're probably priced way too high for the average tech (US\$1,500 and higher). Second, you can use the method I use—*replace and pray*. Open the system case and replace each stick, one at a time, with a known good replacement stick. (You have one of those lying around, don't you?) This method, although potentially time-consuming, certainly works. With PC prices as low as they are now, you could simply replace the whole system for less than the price of a dedicated RAM tester.

Third, you could run a software-based tester on the RAM. Because you have to load a software tester into the memory it's about to scan, there's always a small chance that simply starting the software RAM tester might cause an error. Still, you can find some pretty good free ones out there. My favorite is the venerable Memtest86 written by Mr. Chris Brady ([www.memtest86.com](http://www.memtest86.com)). Memtest86 will exhaustively check your RAM and report bad RAM when it finds it (Figure 4-37).

```

Memtest86+ v1.65 : Pass 11% ####
Athlon 64 (0.09) 2009 MHz : Test 12% ####
L1 Cache: 128K 16468MB/s : Test #4 [Moving inversions, random pattern]
L2 Cache: 512K 16468MB/s : Testing: 100K - 256M 256M
Memory : 256M 11224MB/s : Pattern: 94a989c0
Chipset : Intel i440BX

-----
WallTime   Cached   RsvdMem   MemMap   Cache   ECC   Test   Pass   Errors   ECC   Errs
-----
0:00:05    256M    216K    e820-Std   on   off   Std    0     0
-----

(ESC)Reboot (c)configuration (SP)scroll_lock (CR)scroll_unlock

```

Figure 4-37 Memtest86 in action

## Beyond A+

### The Next Generations

Computer games have done more for advances in PC hardware than anything else. Have you ever seen a late-generation game such as Half-Life 2 or Far Cry? These games require powerful video cards, and every video card comes with DRAM. To make these beautiful games, the video card's DRAM needs to be incredibly fast. Video card makers adopted many DRAM technologies such as DDR even before they were popular as system RAM on PCs. If you want to see where DRAM is going, check out video cards (Figure 4-38).



**NOTE** Go to Chapter 17 for a more detailed description of video memory.



**Figure 4-38** Half-Life 2—these beautiful graphics are possible (in part) because of fast DRAM.

Two advancements to DDR2 originally found only in video cards will show up as primary system memory within a year of this writing: DDR3 and DDR4. DDR3 uses even less power than DDR2 and doubles again the speed of the DRAM's I/O. DDR4 chips push the speed even higher than DDR3!

DDR3 and DDR4 will almost certainly appear in the upcoming *fully buffered DIMMs (FB-DIMMs)*. FB-DIMMs are registered—but with a twist. They replace the 64-bit-wide DIMM data connection with a staggeringly fast serial connection, expected to reach as much as 4.8 GBps! That should be fast enough to keep our PCs happy for the next few years!

## Chapter Review Questions

1. Steve adds a second 512-MB 240-pin DIMM to his PC, which should bring the total RAM in the system up to 1 GB. The PC has a Pentium 4 3.2-GHz processor and three 240-pin DIMM slots on the motherboard. When he turns on the PC, however, the RAM count only shows 512 MB of RAM. Which of the following is most likely to be the problem?
  - A. Steve failed to seat the RAM properly.
  - B. Steve put DDR SDRAM in a DDR 2 slot.

- C. The CPU cannot handle 1 GB of RAM.
  - D. The motherboard can use only one RAM slot at a time.
2. Scott wants to add 512 MB of PC100 SDRAM to an aging but still useful desktop system. The system has a 100-MHz motherboard and currently has 256 MB of non-ECC SDRAM in the system. What else does he need to know before installing?
- A. What speed of RAM does he need?
  - B. What type of RAM does he need?
  - C. How many pins does the RAM have?
  - D. Can the system handle that much RAM?
3. What is the primary reason that DDR2 RAM is potentially faster than DDR RAM?
- A. The core speed of the RAM chips is faster.
  - B. The input/output speed of the RAM is faster.
  - C. DDR RAM is single-channel and DDR2 RAM is dual-channel.
  - D. DDR RAM uses 184-pin DIMMs and DDR2 uses 240-pin DIMMs.
4. What is the term for the delay in the RAM's response to a request from the MCC?
- A. Variance
  - B. MCC gap
  - C. Latency
  - D. Fetch interval
5. Rico has a motherboard with four RAM slots that doesn't seem to work. He has two RDRAM RIMMs installed, for a total of 1 GB of memory, but the system won't boot. What is likely to be the problem?
- A. The motherboard requires SDRAM, not RDRAM.
  - B. The motherboard requires DDR SDRAM, not RDRAM.
  - C. The motherboard requires all four slots filled with RDRAM.
  - D. The motherboard requires that the two empty slots be filled with CRIMMs for termination.
6. Silas has a new Athlon 64 motherboard with two sticks of DDR2 RAM installed in two of the three RAM slots, for a total of 2 GB of system memory. When he runs CPU-Z to test the system, he notices that the software claims he's running single-channel memory. What could be the problem? (Select the best answer.)
- A. His motherboard only supports single-channel memory.
  - B. His motherboard only supports dual-channel memory with DDR RAM, not DDR2.

- C. He needs to install a third RAM stick to enable dual-channel memory.
  - D. He needs to move one of the installed sticks to a different slot to activate dual-channel memory.
7. Motherboards that support more than four sticks of RAM may require what to function properly?
- A. Buffered RAM
  - B. ECC RAM
  - C. Dual-channel RAM
  - D. DDR2 RAM
8. What is the best way to determine the total capacity and specific type of RAM your system can handle?
- A. Check the motherboard book.
  - B. Open the case and inspect the RAM.
  - C. Check the Device Manager.
  - D. Check the System utility in the Control Panel.
9. Gregor installed a third stick of known-good RAM into his Athlon 64 system, bringing the total amount of RAM up to 768 MB. Within a few days, though, he started having random lockups and reboots, especially when doing memory intensive tasks like gaming. What is most likely the problem?
- A. Gregor installed DDR RAM into a DDR2 system.
  - B. Gregor installed DDR2 RAM into a DDR system.
  - C. Gregor installed RAM that didn't match the speed or quality of the RAM in the system.
  - D. Gregor installed RAM that exceeded the speed of the RAM in the system.
10. Cindy installs a second stick of DDR2 RAM into her Core 2 Duo system, bringing the total system memory up to 2 GB. Within a short period of time, though, she begins experiencing Blue Screens of Death. What could the problem be?
- A. She installed faulty RAM.
  - B. The motherboard could only handle 1 GB of RAM.
  - C. The motherboard needed dual-channel RAM.
  - D. There is no problem. Windows always does this initially, but gets better after crashing a few times.

---

## Answers

1. A. Steve failed to seat the RAM properly.
2. D. Scott needs to know if the system can handle that much RAM.
3. B. The input/output speed of DDR2 RAM is faster than DDR RAM (although the latency is higher).
4. C. Latency is the term for the delay in the RAM's response to a request from the MCC.
5. D. RDRAM-based motherboards require empty slots to be filled with CRIMMs for termination.
6. D. Motherboards can be tricky and require you to install RAM in the proper slots to enable dual-channel memory access. In this case, Silas should move one of the installed sticks to a different slot to activate dual-channel memory. (And he should check the motherboard manual for the proper slots.)
7. A. Motherboards that support more than four sticks of RAM may require buffered RAM to function properly.
8. A. The best way to determine the total capacity and specific type of RAM your system can handle is to check the motherboard book.
9. C. Most likely, Gregor installed RAM that didn't match the speed or quality of the RAM in the system.
10. A. If you have no problems with a system and then experience problems after installing something new, then chances are the something new is at fault!