CHAPTER 3

Understanding RAM and CPUs

In this chapter, you'll learn about two important concepts for any A+ technician to understand: random access memory (RAM) and central processing units (CPUs). A CPU is the brain of the computer, performing most of the processing, and RAM is used to store applications and data being used by the CPU. Both continue to be steadily improved and include a significant amount of technical detail that can easily confuse a regular user. This chapter will help you understand many of the terms used when describing them.

Exam 220-801 objectives in this chapter:

- 1.2 Differentiate between motherboard components, their purposes, and properties.
  - CPU sockets
- 1.3 Compare and contrast RAM types and features.
  - Types
    - DDR
    - DDR2
    - DDR3
    - SDRAM
    - SODIMM
    - RAMBUS
    - DIMM
    - Parity vs. non-parity
    - ECC vs. non-ECC
    - RAM configurations
      - Single channel vs. dual channel vs. triple channel
    - Single sided vs. double sided
  - RAM compatibility and speed

- 1.6 Differentiate among various CPU types and features and select the appropriate cooling method.
  - Socket types
    - Intel: LGA, 775, 1155, 1156, 1366
    - AMD: 940, AM2, AM2+, AM3, AM3+, FM1, F
  - Characteristics
    - Speeds
    - Cores
    - Cache size/type
    - Hyperthreading
    - Virtualization support
    - Architecture (32-bit vs. 64-bit)
    - Integrated GPU
  - Cooling
    - Heat sink
    - Fans
    - Thermal paste
    - Liquid-based

Exam 220-802 objectives in this chapter:

- 4.2 Given a scenario, troubleshoot common problems related to motherboards, RAM, CPU and power with appropriate tools.
  - Common symptoms
    - Unexpected shutdowns
    - System lockups
    - Overheating

RAM

When technicians are talking about a computer's memory, they are primarily talking about *random access memory (RAM)*. RAM is used for short-term storage of applications or data so that the processor can access and use this information. In contrast, computers use hard drives for long-term storage of data.

Most RAM is volatile. This doesn't mean that it's explosive; it means that data in RAM is lost when power is removed.

As an introduction, the following list identifies commonly used types of RAM. All of these types of RAM are volatile.

- **Dynamic RAM (DRAM).** Dynamic refers to how bits are stored in an electrical component called a capacitor. The capacitor holds the bit as a charge, but the capacitor needs to be regularly refreshed to hold the charge. This configuration uses very few components per bit, keeping the cost low, but the constant refresh reduces the speed.
- **Synchronous DRAM (SDRAM).** *SDRAM* is synchronized with a clock for faster speeds. Almost all primary *DRAM* used in computers today is SDRAM, but it's often listed as DRAM to avoid confusion with *SRAM*.
- **Static RAM (SRAM).** Static RAM uses switching circuitry instead of capacitors and can hold a charge without a constant refresh. It requires more components per bit so it is more expensive, but due to how the switching works, it is quicker than DRAM. Due to the speed, SRAM is commonly used for CPU cache (described later in this chapter) but is rarely used as the primary RAM because of its cost.

NOTE **SRAM VS. SDRAM**

SRAM and SDRAM are often conflated; however, they are different, and the S makes the difference. The *S* in SRAM indicates *static*, but the *S* in SDRAM indicates *synchronous*. Because of its speed, SRAM is used for CPU cache. SDRAM is used as the primary RAM in computer (PCs). Almost all DRAM in personal computers is SDRAM.

Flash memory is very popular, but not as the primary RAM used in a system. USB flash drives, solid-state drives (SSDs), and memory cards used in cameras and other mobile devices all use flash memory. Flash memory is used for BIOS in many motherboards. Unlike DRAM and SRAM, flash memory is not volatile and retains data without power.

Double Data Rate SDRAM

While the original SDRAM versions were quick and efficient for their time, manufacturers have steadily improved them. *Double data rate* (*DDR*) is one of the improvements and is used in almost all SDRAM. As a reminder, SDRAM is tied to a clock, and when the clock ticks, data is transferred.

SDRAM uses only the leading edge for the clock. However, each of the *DDR SDRAM* versions uses both the leading and trailing edge of the clock. This is often called *double pumping*. Figure 3-1 compares the two over two cycles of a clock. You can see that SDRAM has two clocks from these cycles and that DDR has four clocks from the same two cycles.

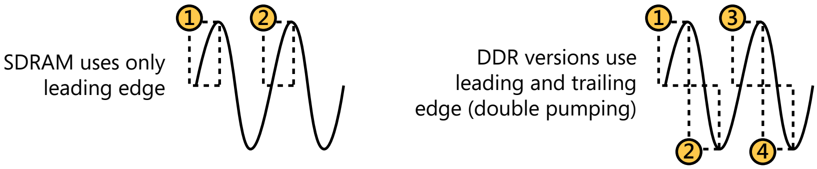

FIGURE 3-1 SDRAM compared with double-pumping DDR.

The following list provides an overview of the different DDR versions:

- **Double Data Rate (DDR) SDRAM.** DDR uses double pumping to double the data rate of SDRAM.
- **DDR2.** DDR2 doubles the data rate of DDR. In addition to double pumping, it modifies the way that data is processed and can transfer twice as much data as DDR SDRAM.
- **DDR3.** DDR3 doubles the data rate of DDR2. It uses double pumping and further modifies the way that data is processed. It can transfer four times as much data as DDR and eight times as much data as SDRAM.

EXAM TIP

DDR3 SDRAM is the primary type of RAM you see in most systems today. It supersedes SDRAM, DDR SDRAM, and DDR2 SDRAM. However, some existing systems have older RAM, and the CompTIA objectives list each type of RAM, so you'll need to be aware of all of them.

DDR4 isn't included in the objectives, but it is on the horizon as a replacement for DDR3. It's expected to double the speed of DDR3.

DIMMs and SODIMMs

RAM comes on cards plugged into the slots in the motherboard. They are smaller than expansion cards, and technicians commonly call memory cards *sticks*. The two most common types of memory sticks are:

- **Dual in-line memory module (DIMM).** A *DIMM* is the circuit board that holds the memory chips.
- **Small outline dual in-line memory module (SODIMM).** *SODIMM* chips are smaller and are used in smaller devices such as laptop computers and some printers.

Figure 3-2 shows a DIMM (top) and a SODIMM (bottom).

IMPORTANT **AVOID ELECTROSTATIC DISCHARGE DAMAGE**

The CPU and RAM are most susceptible to electrostatic discharge (ESD) damage. If you plan on touching the CPU or RAM, ensure that you use ESD wrist straps and other ESD protection as mentioned in Chapter 1, "Introduction to Computers."

FIGURE 3-2 Comparing a DIMM and a SODIMM.

DIMMs and SODIMMs have a different number of pins depending on the type used.

- **DDR SDRAM DIMM:** 184 pins
- **DDR2 SDRAM DIMM:** 240 pins
- **DDR3 SDRAM DIMM:** 240 pins
- **DDR SDRAM SODIMM:** 200 pins
- **DDR2 SDRAM SODIMM:** 144 or 200 pins
- **DDR3 SDRAM SODIMM:** 204 pins

Single Channel, Dual Channel, and Triple Channel

Many motherboards and CPUs support single-channel, dual-channel, and triple-channel memory architectures. Each *single channel* represents a separate 64-bit line of communication that can be accessed independently. With dual channel, the system can access 128 bits at a time; triple channel gives it access to 192 bits at a time.

Using dual and triple channels provides an additional performance enhancement to DDR, DDR2, and DDR3, in addition to double pumping and other enhancements provided by the DDR versions. If you use a dual-channel motherboard with DDR3, it doubles the throughput of DDR3, providing 16 times more data throughput than SDRAM.

If you are upgrading a computer's memory, it's important to understand these channels. You can purchase DIMMs in matched pairs. Where you install each DIMM determines how many channels your system will use and can affect the performance of RAM.

Single Channel vs. Dual Channel

Dual-channel motherboards are very common. If you look at a dual-channel motherboard, you see that it has four memory slots, two slots of one color and two slots of another color. Figure 3-3 shows a diagram of four memory slots labeled for a motherboard using an Intel-based CPU. Slots 1 and 3 are one color, and slots 2 and 4 are another color.

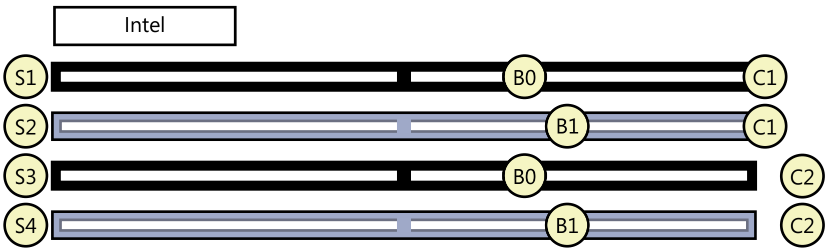

FIGURE 3-3 Intel-based DDR slots (S), banks (B), and channels (C).

- **Slots:** Each slot can accept one DIMM.
- **Banks:** A bank is composed of two slots. In Figure 3-3, Bank 0 includes slots 1 and 3 and these two slots are normally blue. Bank 1 includes slots 2 and 4 and these slots are normally black. This is standard for Intel CPU-based motherboards.
- **Channels:** Each channel represents a separate 64-bit communication path. Slots 1 and 2 make up one channel, and slots 3 and 4 make up the second channel.

EXAM TIP

On most motherboards, the slots are color-coded to identify the banks. Slots of the same color indicate the same bank, and matched pairs should be installed in these slots.

You can install a single DIMM in slot 1, and the system will have a single-channel RAM. You can purchase DIMMs in matched pairs, and it's important to know in which slots to install them. For the best performance, you should install matched DIMMs in the same bank. Looking at Figure 3-3, you should install the matched pair of DIMMs in slots 1 and 3 (Bank 0), leaving slots 2 and 4 empty. The system will take advantage of the dual-channel architecture by using two separate 64-bit channels.

What happens if you install the DIMMs in slots 1 and 2 instead? The system will still work; however, both DIMMs are installed in channel 1, so the system will work with only a single channel. RAM will be about half as fast as it could be if it were installed correctly to take advantage of the dual channels.

Figure 3-3 and the previous explanation describe the color coding, banks, and channels for Intel-based CPU motherboards. However, most motherboards designed for AMD CPUs are organized differently, as shown in Figure 3-4. On these motherboards, slots 1 and 2 make up Bank 0, and slots 3 and 4 make up Bank 1. Channel 1 includes slots 1 and 3, and channel 2 includes slots 2 and 4.

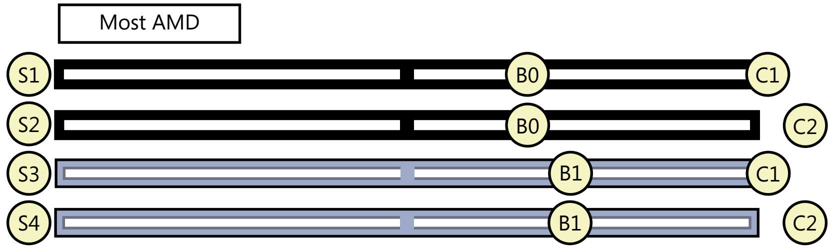

FIGURE 3-4 AMD-based DDR slots (S), banks (B), and channels (C).

While this can be confusing between different motherboards, the good news is that most motherboard manufacturers use the same color for each bank. For Intel-based motherboards, Bank 0 includes slots 1 and 3, and these will be the same color (often blue). Bank 1 includes slots 2 and 4, and they will be a different color (often black). AMD motherboards also use one color for Bank 0 (slots 1 and 2) and another color for Bank 1 (slots 3 and 4).

EXAM TIP

Many motherboards allow the use of different size DIMMs in different channels. However, for the system to use the multichannel capability, each DIMM within a bank must be the same size. If one DIMM in a bank is 1 GB and the second DIMM in the bank is 2 GB, the sizes are different and the system will use single channel. Also, you can use different speed DIMMs in the same bank, although this is not recommended. The speed of the bank will default to the lower-speed DIMM or, in some cases, to single channel.

Triple Channel

On some motherboards, you see six DIMM slots instead of four. This indicates the system supports *triple-channel* memory usage. Table 3-1 shows the configuration of the slots, banks, and channels for a motherboard using triple-channel RAM.

TABLE 3-1 Triple-Channel DIMMs

Slots	Banks	Channels
Slot 1	Bank 0	Channel 1
Slot 2	Bank 1	Channel 1
Slot 3	Bank 0	Channel 2
Slot 4	Bank 1	Channel 2
Slot 5	Bank 0	Channel 3
Slot 6	Bank 1	Channel 3

Slots in each bank are commonly the same color, so you might see a motherboard with Bank 0 slots (slots 1, 3, and 5) all blue and with Bank 1 slots all black.

Triple-channel DIMMs are sold in matched sets of three, similar to how dual-channel DIMMs are sold in matched pairs. When you install triple-channel DIMMs, you should install the matched set in the same bank. For example, if you bought one set, you'd install it in slots 1, 3, and 5.

> ***NOTE*** **QUAD CHANNEL**
>
> **Quad-channel motherboards are also available and have eight DIMM slots. When buying RAM for a quad-channel motherboard, you buy the RAM in a matched set of four. Quad-channel RAM is not mentioned in the CompTIA A+ objectives.**

Single Sided vs. Double Sided

You'd think that *single-sided* and *double-sided* RAM refers to how many sides of a DIMM have chips. That makes sense, but it's not entirely accurate. Instead, single sided or double sided refers to how a system can access the RAM.

In double-sided RAM, the RAM is separated into two groups known as ranks, and the system can access only one rank at a time. If it needs to access the other rank, it needs to switch to the other rank. In contrast, single-sided (or single-rank) RAM is in a single group; the system can access all RAM on the DIMM without switching.

If you have a DIMM with chips on only one side, it is most likely a single-sided (single-rank) DIMM. However, if it has chips on both sides, it can be single rank, dual rank, or even quad rank. You often have to dig into the specs to determine how many ranks it is using.

Usually, you'd think that *double* is better than *single*, but in this case, more rank is not better. Switching back and forth between ranks takes time and slows down the RAM. Single-sided RAM doesn't switch, and if all other factors are the same, single-sided RAM is faster than double-sided RAM.

> ***NOTE*** **DUAL-SIDED IS NOT DUAL CHANNEL**
>
> **Dual-sided (or dual-ranked) is not the same as dual channel. Dual channel improves performance, but a dual-ranked DIMM doesn't perform as well as a single-ranked DIMM.**

Quick Check

1. **A system has six RAM slots. What does this indicate?**
2. **Where should you install two new DIMMs on a dual-channel motherboard?**

Quick Check Answers

1. **Triple-channel RAM.**
2. **In the same bank, identified by slots of the same color.**

RAM Compatibility and Speed

An important point about DDR, DDR2, and DDR3 is that they aren't compatible with each other. You can't use any version in a slot designed for another type. For example, you can use DDR3 DIMMs only in DDR3 slots. From a usability perspective, that's not so great, but if you're trying to remember which types are compatible, it's a lot easier. You can't mix and match them.

Figure 3-5 shows a comparison of the keyings of DDR, DDR2, and DDR3, with a dotted line as a reference through the middle of each one. You can see that the notched key at the bottom of the circuit card is different for each. The standards aren't compatible, and this keying prevents technicians from inserting a DIMM into the wrong slot.

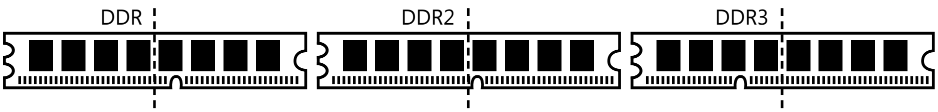

FIGURE 3-5 Comparing DDR versions.

Speeds

Some RAM is faster than other RAM, and with faster RAM you often see faster overall performance. As you'd expect, faster RAM is more expensive. If you're shopping for RAM, you want to ensure that you buy exactly what you need. This includes the correct DDR version, the correct number of channels if your motherboard supports multiple channels, and the correct speed.

The speed of RAM is expressed as the number of bytes it can transfer in a second (B/s) or, more commonly, as megabytes per second (MB/s). However, the speed of most RAM isn't listed plainly. Instead, it's listed using standard names and module names such as DDR3-800 or PC3-12800, respectively. These names indicate their speed, but not directly. If you need to shop for RAM, you need to understand these names and how they relate to the speed.

NOTE **DIMM STICKERS**

Most DIMMs have stickers on them that include the standard name, module name, or both. If you're working at a company that has stored excess DIMMs in static-free packaging, you can identify details from these names.

You can calculate the overall speed of any SDRAM DDR type by using a specific mathematical formula for that type. The formula includes the speed of the clock (Clk), a clock multiplier (Clk Mult) for DDR2 and DDR3, and doubling from double pumping (DP). The speed is calculated for a single channel, which is 64 bits wide, and then converted to bytes by dividing it by 8. The following formulas show how to calculate the speed of each of the DDR versions by using a 100-MHz clock:

- **DDR speed calculation:**
  - Clk × 2 (DP) × 64 (bits) / 8 (bytes)
  - 100 MHz × 2 × 64 / 8 = 1,600 MB/s
- **DDR2 speed calculation:**
  - Clk × 2 (Clk Mult) × 2 (DP) × 64 (bits) / 8 (bytes)
  - 100 MHz × 2 × 2 × 64 / 8 = 3,200 MB/s
- **DDR3 speed calculation:**
  - Clk × 4 (Clk Mult) × 2 (DP) × 64 (bits) / 8 (bytes)
  - 100 MHz × 4 × 2 × 64 / 8 = 6,400 MB/s

Table 3-2 shows how these speeds relate to the different naming conventions used with DDR types. You can see that the standard name is derived from the clock, the clock multiplier, and double pumping. For example, DDR3 uses a 4-times multiplier and double pumping. Therefore, it's eight times faster than SDRAM. The standard name is derived by multiplying the clock by 8. The module name is a little more cryptic, but if you calculate the speed by using the clock, you can see that the PC name indicates the calculated speed in MB/s. Also, you can see that the names include the version (DDR, DDR2, or DDR3).

TABLE 3-2 DDR Standard Names and Module Names

	100 MHz	**166 2/3 MHz**	**200 MHz**
DDR Standard Name DDR Module Name	DDR-200 PC-1600	DDR-333 PC-2700	DDR-400 PC-3200
DDR2 Standard Name DDR2 Module Name	DDR2-400 PC2-3200	DDR2-667 PC2-5300 PC2-5400	DDR2-800 PC2-6400
DDR3 Standard Name DDR3 Module Name	DDR3-800 PC3-6400	DDR3-1333 PC3-10600	DDR3-1600 PC3-12800

> ***NOTE*** **SOME ROUNDING ALLOWED**
>
> **If you enjoy math, you can plug the fractional number 166 2/3 into the speed calculation formulas and see that they don't work out exactly. For example, DDR2-667 works out to about 5333.312 MB/s. Some manufacturers advertise this as PC2-5300, while others round it up to PC2-5400.**

Each DDR version supports multiple clock speeds, and each newer version supports faster clocks. Some of the clock speeds supported by different DDR versions are as follows:

- **DDR:** 100, 133 1/3, 166 2/3, and 200 MHz
- **DDR2:** 100, 133 1/3, 166 2/3, 200, and 266 2/3 MHz
- **DDR3:** 100, 133 1/3, 166 2/3, 200, 266 2/3, and 400 MHz

A key consideration when purchasing RAM is to ensure that the RAM speeds are supported by the motherboard. If the speeds don't match, the motherboard defaults to the slower speed. For example, if your motherboard has a 100-MHz clock and you install PC3-12800 RAM, the RAM will run at 100 MHz instead of 200 MHz. It still works, but you won't get the benefit of the higher-speed RAM.

EXAM TIP

You might need to shop for memory, either to replace memory in your own system or to help someone else. If you can master how memory is named and marketed, you'll be able to identify the correct memory to purchase.

Compatibility within Banks

In addition to matching the RAM speed with the motherboard speed, you should also match the RAM speed within banks when using dual-channel and triple-channel configurations. If one DIMM in a bank fails, you should replace both with a matched set. However, if you have to replace the failed DIMM with a spare, look for a spare that uses the same speed.

For example, if Bank 0 currently has two PC3-12800 sticks and one fails, you should replace the failed stick with a PC3-12800 stick. PC3-12800 uses a 200-MHz clock. If you replaced it with a PC3-6400 (designed for a 100-MHz clock), both sticks would run at the slower speed or revert to single channel.

REAL WORLD **USING THE WRONG SLOTS RESULTS IN SLOWER RAM**

I once helped a friend troubleshoot the speed of a PC after a RAM upgrade. The system started with two 2-GB RAM sticks installed in slots 1 and 2, incorrectly using a single-channel configuration. These DIMMs were PC3-6400, using a 100-MHz clock, and they were working fine, but he wanted more RAM.

He purchased two new 2-GB PC3-12800 DIMMs designed to work with a 200-MHz clock. His motherboard supported 200 MHz, so it could take advantage of the faster RAM. However, after installing the RAM, he ran some tests and found that all the DIMMs were using 100 MHz, so he called me for some help.

Do you see the problem? It took a while to figure out and was exacerbated by the original RAM using the wrong slots. Bank 0 (in slots 1 and 3) now included one 100-MHz DIMM and one new 200-MHz DIMM, so it ran at the slower speed of 100 MHz. Similarly, Bank 1 (in slots 2 and 4) now included one 100-MHz DIMM and one 200-MHz DIMM, so it also ran at the slower speed.

Most users won't test the speed of the RAM after installing it. They're just happy that they have more memory. However, when speeds are mixed in the same bank, users won't get the higher performance.

Shopping for RAM

When shopping for RAM, you need to determine the clock speed of your computer and then determine the DDR name. You can boot into BIOS, as shown in Chapter 2, "Understanding Motherboards and BIOS," to identify the clock speed used by RAM and then plug it into the formula to determine the standard name and module name.

If you have access to the Internet, there's an easier way. You can go to one of the memory sites, such as *Crucial.com* or *Kingston.com*, and use one of their tools. You can enter the make and model of your computer, and the tool will tell you what memory is supported. *Crucial.com* also has an application that you can download and run to identify your motherboard, the type and speeds of supported RAM, how much RAM is installed, and recommendations for upgrading the RAM. Another tool that can help is CPU-Z (described at the end of this chapter).

EXAM TIP

When shopping for memory, you'll find that most memory resellers use the module name, such as PC3-6400. You'll need to match this with the speed of the clock on the target system. Also, remember that the DDR versions are not compatible. PC2-6400 indicates DDR2, and PC3-6400 indicates DDR3.

Parity and ECC

Desktop systems rarely need extra hardware to detect or correct memory errors, but some advanced servers need this ability. The two primary error-detection technologies are *parity* and *error correction code (ECC)*. When shopping for RAM on desktop systems, you'll almost always buy non-parity and non-ECC RAM.

NOTE **APPLICATIONS CHECK FOR ERRORS**

Applications routinely check for errors and often detect and correct errors without the need for parity or ECC RAM.

Parity works by using 9 bits for every byte instead of 8 bits. It sets the ninth bit to a 0 or a 1 for each byte when writing data to RAM. Parity can be odd parity or even parity, referring to odd and even numbers.

Odd parity is common, and when used, it ensures that the 9 bits always have an odd number of 1s. For example, if the 8 data bits were 1010 1010, it has four 1s. Four is an even number, so the parity bit needs to be a 1. Whenever data is written to RAM, the parity bit is calculated and written with each byte.

When the data is read, the system calculates the parity from the 9 bits. If it ever detects an even number of 1s, it knows there is an error, meaning that the data isn't valid and should not be used. Parity can't fix the problem; it just reports the error.

ECC RAM uses additional circuitry and can detect and correct errors. This extra circuitry adds significantly to the cost of the RAM and should be purchased only when necessary. For example, spacecraft that might be exposed to solar flares commonly use ECC RAM. Additionally, some high-end scientific and financial servers need it to ensure that the data in RAM remains error-free.

Rambus and RDRAM

Another type of DRAM is *Rambus DRAM (RDDRAM)*. More commonly, you see it referred to as Rambus, Rambus DRAM, or RDRAM. RDRAM is not compatible with any of the DDR versions and is rarely used.

The circuit boards are called *Rambus in-line memory modules (RIMMs)* instead of DIMMs. When installing RDRAM, you must install it in pairs. In some cases, only one circuit card has memory and the second circuit card in the pair is needed to complete the circuit. The second card is called a *continuity RIMM (CRIMM)*.

EXAM TIP

Rambus and RDRAM are mentioned in the CompTIA objectives, but don't be surprised if you never see a RIMM. They aren't used in new computers, but you might see one in an older computer. You can identify RIMMs by the distinctive metal covering over the chips.

RDRAM generates quite a bit of heat. To dissipate the heat, the chips are covered with a piece of metal acting as a heat sink or heat spreader. This makes them easy to identify because DDR SDRAM is not covered with metal.

CPUs

The *processor*, or *central processing unit (CPU)*, is the brain of the computer. It does the majority of the processing work and is a key factor in the overall performance of a system. Over the years, CPUs have steadily improved, and as a computer technician, you're expected to know some basics about them.

There are two primary manufacturers of computers used in computers: *Intel* and *Advanced Micro Devices (AMD)*.

- **Intel.** Intel is the largest seller of CPUs, selling about 80 percent to 85 percent of all CPUs. It manufactures other products as well, including chipsets, motherboards, memory, and SSDs.
- **AMD.** AMD is the only significant competition to Intel for CPUs, and it sells about 10 percent to 15 percent of all CPUs. It also manufactures other products, including graphics processors, chipsets, and motherboards.

It's possible to purchase a new CPU and install it in a motherboard as part of an upgrade. An important question to ask is, "What should I buy?" When shopping, you'll see names like the following:

- Intel Core i7-960 Processor 3.2 GHz 8 MB Cache Socket LGA 1366
- Phenom II X4 965 AM3 3.4 GHz 512KB 45 NM

Will either of these fit in your motherboard? You might not know right now, but by the end of this chapter, you'll have the information to answer that question.

> ***NOTE*** **RISC**
>
> **You might hear about Advanced RISC Machine (ARM) processors. ARM uses a *reduced instruction set computer (RISC)* architecture and often runs more quickly and with less power than Intel and AMD-based CPUs, so these processors don't need fans. ARM processors are popular in tablets such as the iPad, but you can't replace CPUs in a tablet. You can replace CPUs in computers, so the Intel and AMD CPUs are more important to understand as a computer technician.**

32-bit vs. 64-bit

CPUs are identified as either 32-bit or 64-bit. Similarly, operating systems and many applications are referred to as either 32-bit or 64-bit. Key points to remember include the following:

- Windows operating systems come in both 32-bit and 64-bit versions.
- A 64-bit CPU is required to run a 64-bit operating system.
- A 64-bit operating system is required for 64-bit applications.
- A 64-bit CPU will also run 32-bit software.

The numbers 32 and 64 refer to the address bus discussed in Chapter 2. As a reminder, the address bus is used to address memory locations. A 32-bit CPU supports a 32-bit address bus and can address 2^{32} memory locations, or 4 GB of RAM. A 64-bit CPU supports a 64-bit address bus and can address 2^{64} memory locations, or about 17 EB.

> ***NOTE*** **NOT REALLY 4 GB**
>
> **The CPU also uses this address bus to address devices in the system in addition to RAM. Because of this, a 32-bit system reserves some of the address space for the other devices. If you install 4 GB of RAM in a 32-bit system, you find that operating system can use only about 3.3 GB.**

Operating systems and applications have gotten more sophisticated over the years. Developers have programmed extra features and capabilities, but all of these extras consume additional RAM. For many users, 4 GB of RAM simply isn't enough.

Due to the demand, developers such as Microsoft have created 64-bit versions of their operating systems. However, these 64-bit operating systems can run only on 64-bit CPUs. If you want to directly address more than 4 GB of RAM, you need both a 64-bit CPU and a 64-bit operating system.

- **32-bit and x86.** You often see 32-bit operating systems and software referred to as x86. This is a reference to the long line of Intel CPUs that ended in 86 and can run 32-bit software. AMD processors have different names but are also known to be x86-compatible.
- **64-bit.** Intel refers to its 64-bit processors as Intel 64, and AMD calls its 64-bit processors AMD64. Software makers often refer to 64-bit compatible software as x64.

EXAM TIP

If you want to use 64-bit operating systems, you must have a 64-bit CPU, but you do not need to have software designed specifically for a CPU model. For example, Windows operating systems will work with either Intel or AMD CPUs.

CPU Cores

Most CPUs today have multiple cores within them. Each *core* is a fully functioning processor. With multiple cores, the CPU can divide tasks among each core. The result is a faster system.

Operating systems view the multiple cores as individual CPUs. For example, a single eight-core processor will appear in Task Manager as though it is eight separate processors, as shown in Figure 3-6.

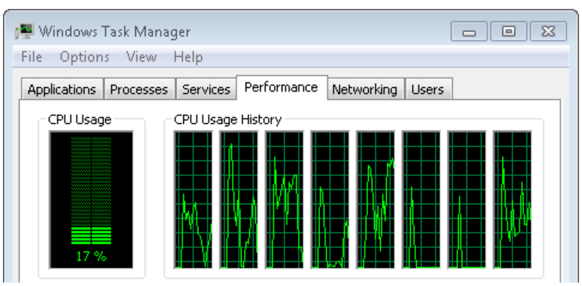

FIGURE 3-6 Task Manager showing eight cores of a single CPU.

MORE INFO **CHAPTER 13, "USING WINDOWS OPERATING SYSTEMS"**

Figure 3-6 shows a partial view of Windows Task Manager. You can start it on Windows systems by pressing Ctrl+Shift+Esc. Chapter 13 provides more details about Task Manager, including how to interpret the displays.

It's worth noting that Figure 3-6 is the same view you'd see if you had an Intel four-core processor with hyper-threading enabled. Hyper-threading is described later in this chapter.

A key point to remember is that even when a CPU has multiple cores, it is still a single chip that plugs into the motherboard. Motherboards are available that accept multiple CPUs, but they are more common on servers than on desktop systems. Most desktop systems have a single CPU, and it's common to see CPUs with multiple cores.

Hyper-Threading

Hyper-Threading Technology (HT) is used on some Intel CPUs to double the number of instruction sets the CPU can process at a time. Within a CPU, a thread is an ordered group of instructions that produce a result. When hyper-threading is used, a single CPU can process two threads at a time.

This is not physically the same as a multiple-core CPU. However, just as a dual-core CPU simulates two physical CPUs, a single-core CPU with hyper-threading simulates two physical CPUs. Operating systems can't tell the difference.

> ***NOTE*** **ENABLE IN BIOS**
>
> **Hyper-threading needs to be enabled in the BIOS before the operating system is installed for it to work. This is usually listed as hyper-threading within a CPU Technology Support menu.**

Intel makes use of both hyper-threading and multiple cores on some of its CPUs. For example, Figure 3-7 shows a screen shot of the System Information tool in Windows 7. It identifies the processor as an Intel Core i7 CPU with four cores and eight logical processors. Each core is using hyper-threading, and the operating system interprets it as eight CPUs.

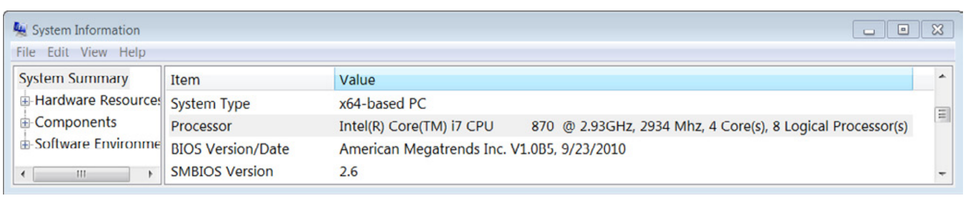

FIGURE 3-7 Msinfo32 showing that hyper-threading is enabled.

> ***MORE INFO*** **CHAPTER 2 AND CHAPTER 14**
>
> **Chapter 2 introduced the System Information tool as a way to check your BIOS version. There are several ways to launch this tool, including entering msinfo32 at the command prompt. Chapter 14, "Using the Command Prompt," covers how to start and use the command prompt.**

CPU Cache

Many computer components and software applications use some type of *cache*. As a simple example, web browsers use a browser cache. When you go to a website, information is transmitted over the Internet and displayed in your web browser, and it is also stored in the browser cache. If you go to the website again, data can be retrieved from the browser cache rather than downloaded from the Internet again. The browser uses different techniques to ensure that it displays current data, but if that data is on your drive, it is displayed much more quickly than it would be if it had to be downloaded again.

The CPU has cache that it uses for fast access to data. If the CPU expects to use some type of information again, it keeps that information in cache. A significant difference between the web browser cache and the CPU cache is that the CPU cache is RAM and the web browser cache is stored as a file on a hard drive.

> ***NOTE*** **CACHE**
>
> **Cache is commonly referred to as an area where data is stored for a short time for easy retrieval. It's important to realize that cache can be memory areas that are volatile or can be temporary files stored on hard drives that are kept after a system is powered down.**

CPU Cache Types

The two primary types of cache used by CPUs are:

- **L1 cache.** This is the fastest, and it's located closest to the CPU. A multiple-core CPU has a separate *L1 cache* located on each CPU core.
- **L2 cache.** *L2 cache* is a little slower than L1 cache, and it is shared by all cores of the CPU. In older systems, L2 cache was stored on the motherboard, but today it is much more common for L2 cache to be part of the CPU.

> ***NOTE*** **L3 CACHE**
>
> ***L3 cache*** **is used on some systems, but it isn't as common as L1 and L2. When used, it can be on the motherboard or on the CPU. It is slower than L2 cache and is shared among all cores.**

Figure 3-8 shows the relationship of the CPUs to cache and RAM installed on the motherboard. In the diagram, the CPU is a two-core CPU, and you can see that the L1 cache is included on each core and that L2 cache is shared by each of the cores. When the CPU needs data, it will check the L1 cache first, the L2 cache next, and then the L3 cache if it exists. If the data isn't in cache, the CPU retrieves it from RAM.

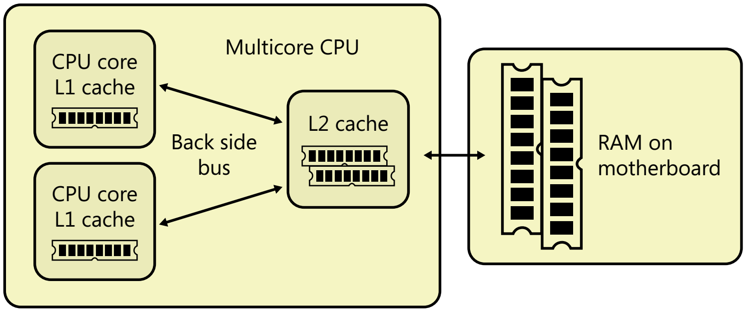

FIGURE 3-8 CPU and cache.

> ***NOTE*** **ACCESSING RAM WITHOUT NORTH BRIDGE.**
>
> **As discussed in Chapter 2, newer CPUs access the motherboard RAM directly, as shown in Figure 3-8. On older CPUs, they access RAM through the north bridge portion of the chipset.**

Many newer CPUs include L1 cache for each core, L2 cache for each core, and a single shared L3 cache—all on the same CPU chip.

Without cache, the CPU would have to store data in the motherboard RAM. The CPU cache is SRAM, which is much faster than the dynamic RAM used on the motherboard. Also, the motherboard RAM is physically farther away, adding more delays.

CPU Cache Size

The size of the CPU cache is small compared to the overall amount of memory in a system. For example, you might see cache sizes as low as 8 KB or as large as 20 MB. In contrast, most personal computers have 1 GB of RAM or more. The cache can be listed as just a total of all L1, L2, or L3 cache, or you might see it listed individually.

- **L1 is smallest.** L1 is sometimes stated as two numbers, such as 32 KB + 32 KB, to indicate it is using one cache for frequently used instructions and another cache for data. Sizes of 32 KB or 64 KB are common.
- **L2 is larger than L1.** When a CPU has separate L2 cache for each core, it is often identified as the amount per core. For example, a two-core CPU with 4 MB total L2 cache can be expressed as 2 × 2 MB, or just 2 MB per core. Sizes of 256 KB, 512 KB, and 1,024 KB are common.
- **L3 is larger than L2.** Sizes between 2 MB and 8 MB are common.

Quick Check

1. Which is faster: L1 or L2?
2. Where is hyper-threading enabled?

Quick Check Answers

1. L1 is the cache closest to the CPU, and it is the fastest.
2. In the BIOS.

Speeds

The speed of a CPU is based on the speed of the crystal and the multiplier. For example, if the crystal speed is 100 MHz and the multiplier is 20, the CPU has a speed of 2 GHz (20 × 100). The faster the speed, the faster the CPU.

You commonly see the speed of the processor listed as only the multiplied speed. For example, in Figure 3-7 you can see that the processor is an Intel Core 7 CPU 870 and the clock is listed as 2.93 GHz. The system is using a 133.333-MHz clock (commonly listed as 133 MHz) and a 22-times multiplier.

NOTE **SPEEDS ARE VARIABLE**

Most current processors can dynamically adjust the speed based on requests from the operating system or an application. When a boost in a CPU core is needed, the operating system can send a signal to make the core run faster. Intel refers to this as Turbo Boost, and AMD refers to it as Turbo Core.

Processors are rated based on the maximum speed they can handle, and more expensive processors can handle faster speeds. You can increase the speed by increasing the clock frequency, increasing the multiplier, or both. Most motherboards have this preselected, but it is sometimes possible to manipulate the clock or the multiplier to overclock the system. In some systems, the BIOS includes a Cell menu that enables you to increase the base frequency and increase the CPU Ratio (multiplier).

EXAM TIP

Overclocking a system is not recommended, but it is frequently done. If you overclock a system, you need to take extra steps to keep it cool, such as using liquid cooling. Liquid cooling is discussed later in this chapter.

Chapter 2 mentions the front side bus (FSB) and how it provides a direct connection between the CPU and the north bridge portion of the chipset. In the past, CPU speeds were stated as the FSB speed. Today, many CPUs have taken over the functionality of the north bridge. The CPU still needs to communicate with the chipset, and there are a few different ways this is done, including the following:

- **Intel Direct Media Interface (DMI).** The DMI can use multiple lanes, similar to Peripheral Component Interconnect Express (PCIe).
- **Intel's QuickPath Interconnect (QPI).** Each core in a processor has a separate two-way 20-lane QPI link to the chipset.
- **HyperTransport.** AMD uses *HyperTransport* with the FSB to increase the speed.

You still see CPUs advertised with a speed that you can use for comparisons. For example, one CPU might have a speed of 2.8 GHz and another might have a speed of 3.4 GHz. It's safe to assume that the 3.4-GHz CPU is faster, but the speed isn't always tied to the FSB.

Virtualization Support

Chapter 2 introduced virtualization concepts and instructions on how to enable virtualization in BIOS. As a reminder, virtualization software allows you to run multiple virtual machines (VMs) as guests within a single physical host computer. The CPU needs to support virtualization, and it usually needs to be enabled in BIOS. On many AMD-based systems, virtualization is enabled by default and cannot be disabled.

Most Intel and AMD CPUs include native support for virtualization. The exception is laptop computers, which sometimes include CPUs that do not support it. Intel refers to its virtualization support as VT-x, and AMD calls its support AMD-V. If you want to verify that a CPU or motherboard supports virtualization, look for those terms.

NOTE **COLD BOOT REQUIRED**

If you change the virtualization setting in the BIOS, it's recommended that you do a cold boot. A cold boot completely powers down the computer. You should wait about 10 seconds and then restart the computer. In contrast, a warm boot shuts down the software and restarts it, but does not shut down the power.

Integrated GPU

Graphics is one of the areas of a computer that has been increasing as quickly as the CPU area, and the two are starting to merge. Early computers could display only letters on a screen 80 characters wide. Today, it's common to watch high-quality video streaming from a website or to play games with computer-generated graphics and amazingly realistic scenery.

The following list describes the progression of graphics capabilities on computers:

- **Onboard graphics.** Graphics capability was built into the chipset. This was often very basic but met most needs.
- **Expansion cards.** You could install a graphics card with a dedicated graphics processing unit (GPU) and plug it into an available expansion slot. Instead of the CPU doing the graphics calculations, the GPU would do them. Peripheral Component Interconnect (PCI) cards were an early version.
- **Dedicated graphics slots.** Accelerated Graphics Port (AGP) provided a single dedicated graphics slot that worked separately from PCI. AGP did not compete with PCI, so it provided better performance. Later, PCIe allowed graphics cards to use their own dedicated lanes, and it replaced AGP.
- **Direct access graphics.** The CPU interacted with the AGP slot via the chipset. Newer CPUs bypass the chipset and interact directly with a dedicated PCIe slot used for graphics. This is common in many systems today.
- **Integrated graphics processing unit (GPU).** A recent trend in newer CPUs is to include an *integrated GPU* on the CPU. GPUs can provide high-quality graphics without the additional cost of a graphics card. However, these are not as powerful as a dedicated card.

AMD refers to some chips with a GPU as an accelerated processing unit (APU) instead of a CPU. APUs can include a GPU or other specialized capability, and the AMD Fusion is an example.

CPU Versions

There is a dizzying number of different processors. You're not expected to know the characteristics of each individual CPU, but you should be able to recognize the names and know the manufacturers. The objectives specifically list the CPU socket types you should know, but for the sockets to make sense, you need to have a little bit of knowledge about the CPU versions.

Intel and AMD use code names related to the manufacturing process and then create different processor families with the process. The manufacturing process is stated as a measurement and refers to the distance between certain components within the chip. Many current CPUs have processes of 65 nanometers (nm), 45 nm, 32 nm, and 22 nm. A nanometer is one billionth of a meter and is often used to express atomic scale dimensions, such as the width of an atom or the width of a group of molecules. In this case, smaller is better.

NOTE **MOORE'S LAW**

One of the founders of Intel, Gordon Moore, predicted in 1965 that the number of transistors that could be placed on a chip would double about every two years. This miniaturization trend has been consistent since his prediction. With more transistors, chips are faster and more complex, and the process used to create them is smaller.

The following are recent Intel and AMD code names:

- **Intel**
  - Core—65-nm and 45-nm process
  - Nehalem—45-nm process
  - Sandy Bridge—32-nm process
  - Ivy Bridge—22-nm process
- **AMD**
  - K8—65-nm, 90-nm, and 130-nm processes
  - K9—processors were never released
  - K10—65-nm process
  - K10.5—45-nm process
  - Bulldozer—22-nm process

Table 3-3 shows a list of common Intel code names and some of their related CPUs. You can see that the Core i3, i5, and i7 family names are frequently repeated.

TABLE 3-3 Intel Code Names and Processors

Architecture Name	CPU Family names
Core	Core 2 Duo, Core 2 Quad, Core 2 Extreme
Nehalem	Intel Pentium, Core i3, Core i5, Core i7, Xeon
Sandy Bridge	Celeron, Pentium, Core i3, Core i5, Core i7
Ivy Bridge	Core i5, Core i7, Xeon

The Core i3, i5, and i7 series represents a Good, Better, Best philosophy, with the i3 versions representing the basic version and the i7 versions providing the most power. The number (such as i3 or i5) doesn't refer to the number of cores.

It's also important to realize that there are significant differences between a Nehalem Core i5 and an Ivy Bridge Core i5. The Ivy Bridge versions have smaller processes and are more powerful.

MORE INFO **WIKIPEDIA**

This chapter does not list all the existing Intel and AMD CPUs. If you want to see a list of Intel or AMD processors, check out these two Wikipedia pages: *http://en.wikipedia.org/wiki/List_of_Intel_microprocessors* and *http://en.wikipedia.org/wiki/List_of_AMD_microprocessors*.

Table 3-4 shows a list of common AMD code names and their related CPUs. The primary AMD CPUs that you find in desktop computers are Sempron, Athlon, and Phenom.

TABLE 3-4 AMD Code Names and Processors

Architecture Name	CPU Family names
K8	Opteron, Athlon 64, Athlon 64 FX, Athlon 64 X2, Sempron, Turion 64, Turion 64 X2
K10	Opteron, Phenom, Athlon, Athlon X2, Sempron
K10.5	Phenom II, Athlon II, Sempron, Turion II
Bulldozer	FX (Zambezi), Interlagos Opteron

EXAM TIP

Many AMD processor names give clues as to what they include. If the name includes 64, it is a 64-bit CPU. When the name has an X (such as X2), it indicates how many cores the processor has.

CPU Socket Types

A CPU plugs into a *socket* on the motherboard. There was a time when just about every motherboard had the same socket type, but that certainly isn't the case today. Instead, there are a wide variety of different socket types for different types of CPUs. If you ever need to replace a CPU, it's important to recognize that there are different types of sockets. The following sections talk about some sockets used by Intel and AMD, with information about how they are installed.

Zero Insertion Force

It's important that each of the pins on a CPU has a good connection to the motherboard. In early versions of CPUs, this was accomplished by creating a tight connection between the pins and the socket. This required technicians to use some force to plug the CPU into the socket. Unfortunately, it was easy to bend one or more pins, and bent pins would often break, making the CPU unusable.

Manufacturers came up with a great idea to eliminate the problem—*zero insertion force (ZIF)* sockets. A ZIF socket has a locking lever. You can place a CPU into a socket without any force other than gravity, and after the CPU is in place, you lock the lever to secure it. This lever ensures that the pins are making a solid connection to the motherboard.

Figure 3-9 shows a ZIF socket with the lever raised. The CPU is removed and standing up on the left. You can see that there are some areas on the CPU where there aren't any pins. These provide a key, and they match up to areas on the socket where there aren't any pin holes.

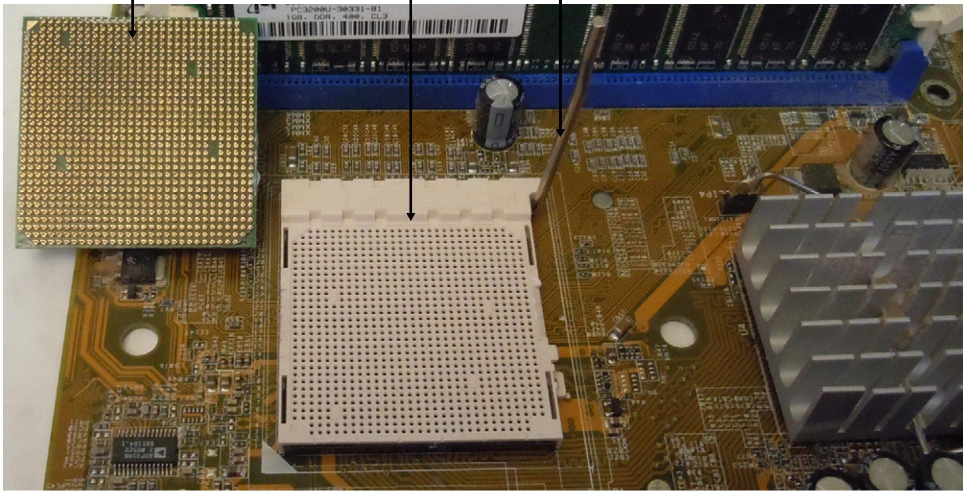

FIGURE 3-9 Processor and ZIF socket.

NOTE **CPUS ARE KEYED**

CPU sockets and CPUs are keyed so that the CPU fits into a socket in only one way. If you feel any resistance when putting a CPU into a ZIF socket, it indicates that the pins aren't lined up. You should double-check the keying and ensure that the CPU is lined up correctly. If you try to force it, you will likely bend some pins and ruin either the chip or the socket.

PGA vs. LGA

The socket shown in Figure 3-9 is a *pin grid array (PGA)* type of socket. It includes holes into which the pins can be plugged. A newer type of socket is a *land grid array (LGA)* socket. Instead of the processor having pins and plugging into a socket with holes, the socket has small pins, and the CPU has small pins created as bumps or pads. When the CPU is installed, the pins and bumps line up, making the connection.

When using an LGA socket, the CPU sits on top of the socket but is locked in place with a flip-top case. Figure 3-10 shows an example of a flip-top case used with an Intel processor.

This socket has a hinged top and a lever that locks the case when it's closed. You unlock the lever, open the case, and remove the CPU. When installing a new CPU, ensure that the keys line up, place the CPU in the case, close the top, and lock it with the lever. Remember to use ESD protection when handling the CPU.

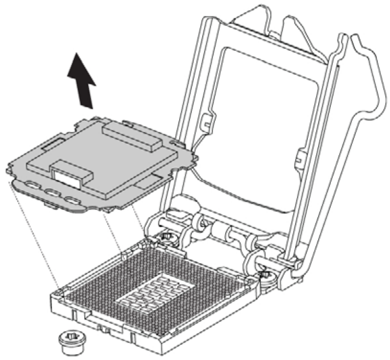

FIGURE 3-10 Removing processor from a flip-top case. Diagram provided by Intel. [Copyright © Intel Corporation. All rights reserved. Used by permission.]

Another type of array you might run across is ball grid array (BGA). In a BGA chip, the pins on the CPU are replaced with balls of solder. The chip is mounted in the socket and then heated, often in an oven, to melt the solder. Manufacturers can fit more pins on a BGA CPU, and they are sometimes used in mobile devices.

Intel CPU Sockets

The following list describes recent Intel sockets:

- **LGA 775.** 775 pins. Also called Socket T. Replaced Socket 478.
- **LGA 1366.** 1,366 pins. Also called Socket B and designed to replace LGA 755 in high-end desktop computers.
- **LGA 2011.** 2,011 pins and released in 2011. Also called Socket R. It replaces LGA 1366 sockets in high-end desktop systems.
- **LGA 1156.** 1,156 pins. Also called Socket H or Socket H1.
- **LGA 1155.** 1,155 pins. Also called Socket H2 and replaces LGA 1156 in basic desktop systems. LGA 1,156 CPUs will work in LGA 1155, but the BIOS may need to be upgraded.

EXAM TIP

Notice that the numbers indicate the number of pins and are not a reflection of newer or older sockets. Also, each of these Intel sockets is an LGA socket.

Table 3-5 lists the common Intel sockets along with some CPUs used with them, busses they support, and supported DDR channels.

TABLE 3-5 Intel Sockets and Related CPUs

Type	CPUs, Busses, DDR Channels
LGA 775 (Socket T)	Pentium 4, Pentium D, Core 2 Duo, Core 2 Quad, Celeron, Xeon Front side bus, single channel DDR2 and DDR3 RAM
LGA 1366 (Socket B)	Core i7, Xeon, Celeron QPI, triple channel DDR3 RAM
LGA 2011 (Socket R)	Core i7, Xeon QPI, DMI, quad channel DDR3 RAM
LGA 1156 (Socket H or H1)	Core i3, Core i5, Core i7, Celeron, Pentium, Xeon DMI, dual channel DDR3 RAM
LGA 1155 (Socket H2)	Core i3, Core i5, Core i7, Celeron, Pentium DMI, dual channel DDR3 RAM

AMD CPU Sockets

The following list describes recent AMD sockets:

- **Socket 940.** 940 pins (PGA).
- **Socket AM2.** 940 pins (PGA). Not compatible with Socket 940.
- **Socket AM2+.** 940 pins (PGA). Replaces AM2. CPUs that can fit in AM2 can also fit in AM2+.
- **Socket AM3.** 941 pins (PGA). Replaces AM2+. Supports DDR3. CPUs designed for AM3 will also work in AM2+ sockets, but CPUs designed for AM2+ might not work in AM3 sockets.
- **Socket AM3+.** 942 pins (PGA). Replaces AM3. CPUs that can fit in AM3 can also fit in AM3+.
- **Socket FM1.** 905 pins (PGA). Used for accelerated processing units (APUs).
- **Socket F.** 1,207 pins (LGA). Used on servers and replaced by Socket C32 and Socket G34.

Table 3-6 lists the common AMD sockets along with some CPUs used with them, busses they support, and supported DDR channels.

TABLE 3-6 AMD Sockets and Related CPUs

Socket	CPUs, Busses, DDR Channels
940	Opteron and Athlon 64 FX FSB with HyperTransport version 1, single channel DDR2 RAM
AM2	Athlon 64, Athlon 64 X2, Athlon FX, Sempron, Phenom, Opteron FSB with HyperTransport version 2, single channel DDR2 RAM
AM2+	Athlon 64, Athlon 64 X2, Athlon II, Sempron, Phenom, Phenom II, Opteron FSB with HyperTransport version 3, single channel DDR2 RAM

AM3	Phenom II, Athlon II, Sempron, Opteron FSB with HyperTransport version 3, single channel DDR2 and dual channel DDR3 RAM
AM3+	Phenom II, Athlon II, Sempron, Opteron FSB with HyperTransport version 3, dual channel DDR3 RAM
FM1	Fusion and Athlon II APUs FSB with HyperTransport version 3, dual channel DDR3 RAM
F	Opteron, Athlon 64 FX FSB with HyperTransport version 3, single channel DDR2 RAM

Comparing Names

Earlier in this chapter, I listed two CPUs using common marketing names. To tie some of this together, here are the two CPUs with an explanation of the names. I'm hoping these names make a lot more sense at this point.

- **Intel Core i7-960 Processor 3.2 GHz 8 MB Cache Socket LGA 1366.** This name indicates that it is an Intel processor in the Core i7 family with a model number of 960 and a 3.2-GHz multiplied clock. The 8-MB cache phrase refers to the total amount of cache. Last, LGA 1366 indicates the type of socket into which the processor will plug.
- **Phenom II X4 965 AM3 3.4 GHz 512 KB 45 NM.** This indicates that it is an AMD Phenom II processor with a model number of 960. X4 indicates that the processor has four cores, and AM3 indicates the socket type. The 3.4-GHz clock speed is the internal speed of the processor. Cache size is indicated by 512 KB, and in this case, it indicates the L2 cache size for each of the cores. The process is 45 nm.

Cooling

CPUs have millions—and sometimes billions—of miniaturized transistors within them, all connected with extremely small wires. If these transistors or wires get too hot, they can easily break, rendering the CPU useless. Manufacturers spend a lot of time designing these chips, and one of their goals is to keep temperatures within acceptable limits. However, most of the cooling occurs externally.

Heat Sinks, Fans, and Thermal Paste

Common methods of cooling a CPU include using a *heat sink*, a fan, and *thermal paste*. Take a look at Figure 3-11 as you read about how these components work together.

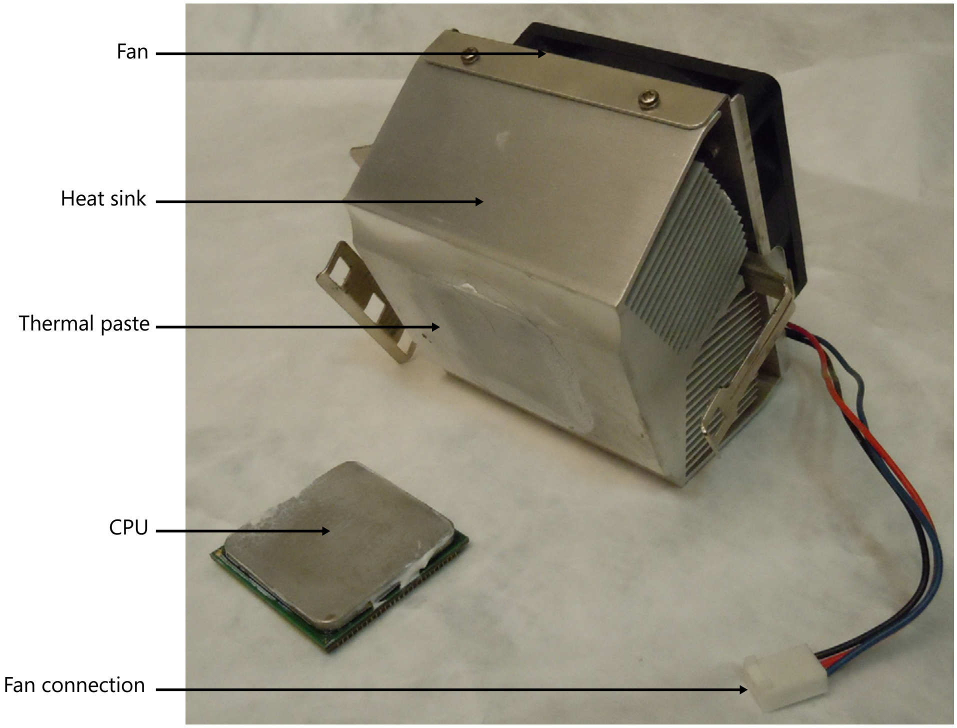

FIGURE 3-11 CPU with heat sink and attached fan.

- **Heat sink.** A heat sink is a piece of metal that draws heat from the CPU and dissipates it into the air. Heat sinks have multiple fins to increase the surface area and to allow air to easily flow through them. The fins are usually flared to allow more air through.
- **Fan.** A fan is attached to the heat sink to increase the airflow around the fins. These are called CPU fans. They aren't attached to the CPU but usually plug into the motherboard close to the CPU. Many CPU fans have variable speeds and spin faster when the CPU gets hotter.
- **Thermal paste.** Heat sinks commonly have clamps to secure them to the motherboard and provide a better connection with the CPU. However, there are microscopic gaps in the metal on both the CPU and the heat sink, so it isn't possible to get 100 percent contact between the components. Thermal paste is used to improve this connection. This paste fills these microscopic gaps and also helps draw heat from the CPU into the heat sink.

EXAM TIP

When replacing a CPU, ensure that you clean off the old thermal paste from the heat sink and apply new thermal paste.

If you are replacing a CPU, you'll need to clean off the old thermal paste from the heat sink. Some vendors sell specialized cleaning compounds to remove old paste, but you can often use cotton swabs and isopropyl alcohol to remove it.

After installing the new CPU into the socket and locking the ZIF arm, place a dab of the paste in the center of the CPU. When you attach the heat sink and clamp it down, the pressure will spread the paste evenly between the heat sink and the CPU. Be careful not to apply too much paste; you need only enough to fill the microscopic gaps between the CPU and the heat sink.

Liquid Cooling

An advanced method of keeping a system cool is using a liquid-based cooling system. Liquid-based cooling systems use water (most commonly) or some other liquid that is pumped through the cooling system.

For example, Figure 3-12 shows a basic diagram of a liquid-based cooling system. A specialized heat sink is attached to the CPU, using thermal paste just like a standard heat sink. However, this heat sink has channels so that the liquid can flow through it. Tubing is connected from the pump to the heat sink, and the pump constantly pumps the liquid through the heat sink.

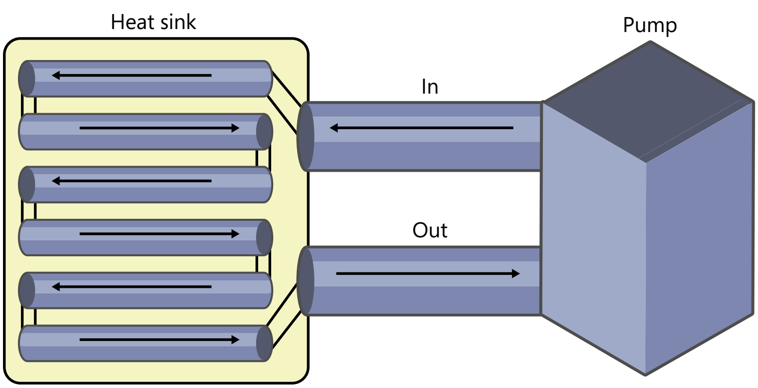

FIGURE 3-12 Liquid-cooled heat sink.

EXAM TIP

Liquid-based cooling can be used for any components that use a heat sink. This includes CPUs, GPUs, and chipsets.

One of the biggest challenges with a liquid-based cooling system is ensuring that the tubing connections do not leak. This is one place where you don't want to skimp on quality. The liquid is usually water, and if it leaks, it could easily destroy the system.

Liquid-based cooling systems are most common among gamers and hobbyists. These people often overclock the processors to get more power out of them, but overclocking generates more heat. Overclocking is sometimes possible by changing jumpers on the motherboard or by manipulating BIOS settings, but manufacturers discourage the practice.

Quick Check

1. **What is another name for Socket H2?**
2. **What's the best way to keep an overclocked CPU cool?**

Quick Check Answers

1. **LGA 1155.**
2. **Liquid cooling.**

Troubleshooting

You might occasionally run across a system that is having a problem with the CPU or RAM. Sometimes the problems are consistent, but more often they are intermittent; sometimes you'll see the problem, sometimes you won't.

Intermittent problems are frequently related to overheating, so a good first step is to ensure that the system has adequate airflow. Shut the system down, open the case, and either vacuum it with an ESD-safe vacuum or take it outside and blow it out with compressed air.

Common Symptoms

The following are some common symptoms and possible causes related to the CPU or RAM:

- **Unexpected shutdowns.** If the system is randomly shutting down or rebooting, the most likely cause is a heat problem. Check the ventilation and clean out the fans.
- **System lockups.** When a computer stops responding to inputs from the keyboard or mouse, technicians refer to it as frozen or locked up. This can also be due to heat issues. Check the ventilation.
- **Continuous reboots.** In some cases, a hardware issue can prevent the system from booting completely. It starts, gets so far, and then resets itself. This is more common after a faulty software update, but it can be due to a hardware problem. If you've just replaced hardware, double-check your steps. If that isn't the issue, boot into Safe Mode and troubleshoot the operating system using the steps provided in Chapter 17, "Troubleshooting Windows Operating Systems."

Tools

If you've cleaned out the system and you're still having intermittent problems, there are two primary things to check:

- **Power supply.** An overloaded or failing power supply can cause intermittent problems. Use a multimeter to verify the voltages. If the voltages are out of tolerance, replace the power supply.
- **RAM.** It is possible to have a certain area of RAM that is faulty. The system can work until it writes data to that area, and then it shuts down or freezes. In some cases, you receive a stop error or blue screen of death (BSOD) with an error code indicating a memory problem. If you suspect a RAM problem, use a memory checker to run memory diagnostics.

EXAM TIP

The two primary *hardware* sources of intermittent problems are the power supply and RAM. The primary *software* source of intermittent problems is a virus or some type of malicious software. Chapter 26, "Recognizing Malware and Other Threats," covers viruses in more depth, but running up-to-date antivirus software usually reveals and removes the problem. Occasionally, you'll need to boot into Safe Mode and run the up-to-date antivirus software.

Windows Memory Diagnostics

Windows Vista and Windows 7 include the Windows Memory Diagnostic tool, and steps later in this section show how to run it. It's easy to run and can perform in-depth testing of the system RAM and the cache within the CPU.

The diagnostics include three sets of tests (basic, standard, and extended). By default, it runs two passes of the standard set of tests, and this is usually good enough. If this passes but you still suspect you have memory problems, you can choose other options by pressing F1 to modify them. For example, if you have an intermittent problem and want to do detailed tests for a day or longer, you can set the pass count to 0 and it will run continuously.

You can use the following steps on a Windows 7 system to run the Windows Memory Diagnostics tool:

1. Click Start and type **Memory** in the Search Programs And Files text box.
2. Select Windows Memory Diagnostic.
3. Select Restart Now and check for problems. After the system reboots, the tests will start and you'll see a display similar to the following graphic. If any errors are identified, they will be displayed in the Status area, but they usually won't stop the diagnostic from running. After the test completes, the system automatically reboots.

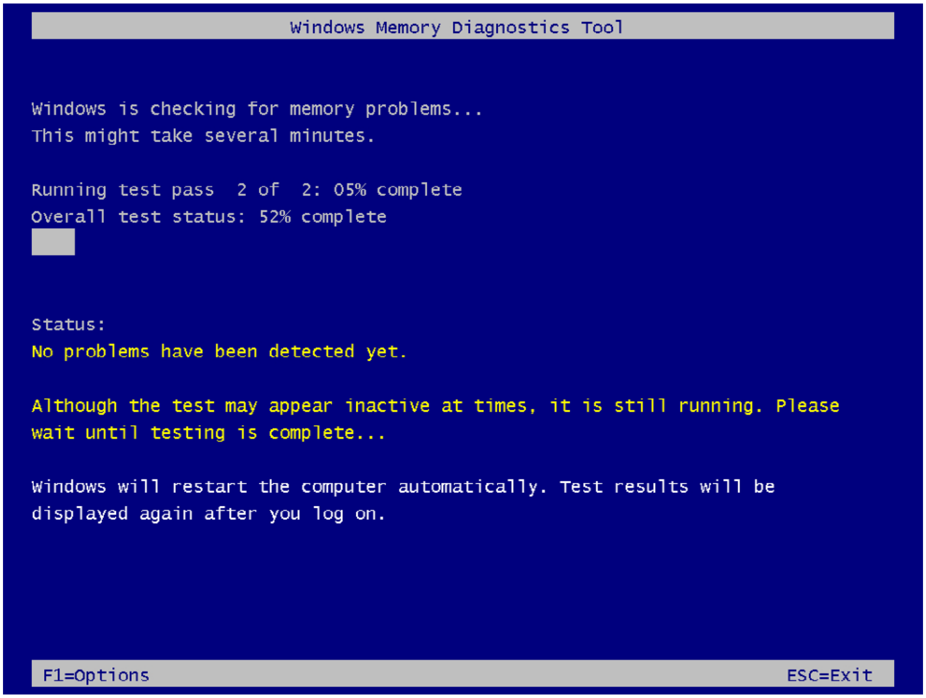

4. About a minute or so after you log on, you'll see a balloon message appear in the system tray at the bottom right indicating the results. It appears and then fades out. If you miss it, you can also view the results in the System log via the Event Viewer. It's listed with a source of MemoryDiagnostics-Results and an Event ID of 1201.

If you're unable to boot into the operating system, you can access the Windows Memory Diagnostic by using several other methods. Each of the following methods will start the Windows Recovery Environment (Windows RE), showing the System Recovery Options, as shown in Figure 3-13. You can then select Windows Memory Diagnostic.

- Press F8 as the system is booting to access the Advanced Boot Options page and select Repair Your Computer.
- Start from a Windows Vista installation DVD, select the Language, and then click Repair Your Computer.
- Create a system repair disc and use it to boot directly into the Windows RE.

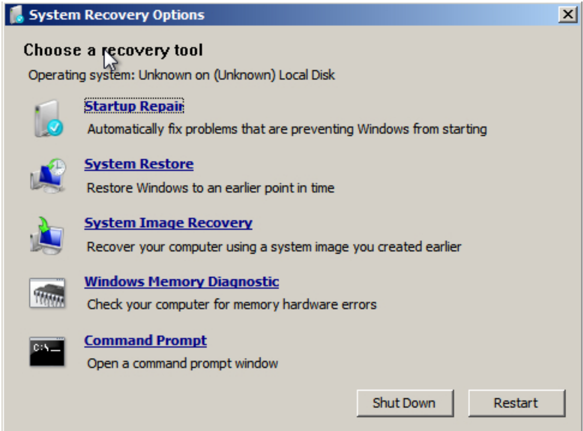

FIGURE 3-13 Running Windows Memory Diagnostic from boot DVD.

MORE INFO **CHAPTER 17**

Chapter 17 covers the Event Viewer, including how to launch it and access different log files. It also includes information on the other system recovery options and how to create a system repair disc in Windows 7.

If the memory diagnostic gives any errors, you might be able to do a quick fix by reseating the memory sticks. Power your system down and open it up. Hook up an ESD strap to ground yourself with the system and then locate the RAM. Press the tabs on each side to pop out each DIMM, and then push each back into the slot until the tabs lock. This same fix can also be used on any expansion card.

You might be wondering why this works. Electrical components expand and contract from heat and cold, causing some movement. Additionally, the electrical contacts can become tarnished, preventing a good connection. When you pop it out and push it back, the friction scrapes the tarnish off the contacts. With the tarnish removed, it has a good connection.

NOTE **CLEANING CONTACTS**

You can clean contacts with contact cleaner created specifically for this purpose. You can also use isopropyl alcohol and a lint-free cloth or cotton swab. You should not rub the contacts with a pencil eraser. The eraser removes the tarnish by scraping it off, but it leaves residue and can cause ESD damage.

CPU-Z

CPU-Z is a handy freeware utility that you can use to view some detailed information on your system. It's been around a long time and has helped many technicians. A copy is on the CD, and you can find a link about the installation here: *getcertifiedgetahead.com/aplus.aspx*.

Figure 3-14 shows a screen shot of the CPU tab of the CPU-Z application. You can see that this provides some detailed information about the processor, clocks, and cache.

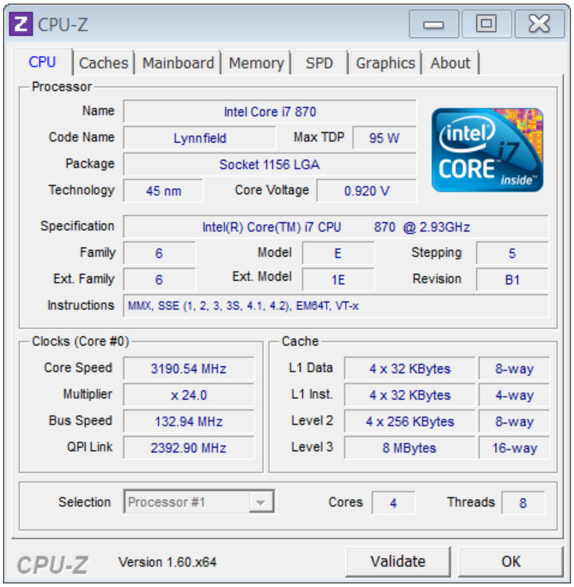

FIGURE 3-14 CPU-Z.

If you click the Mainboard tab, it gives you details about your motherboard and BIOS. The Memory tab provides overall information about installed memory, and the SPD tab enables you to select individual memory slots to determine what is installed. As you'd expect, the Graphics tab provides details about the graphics card. The About tab includes buttons you can use to save the details of the report as either a text file or an HTML file.

Quick Check

1. What are two primary hardware problems that can cause system fails?
2. Where can you determine how much RAM is installed in a system?

Quick Check Answers

1. Faulty power supply or faulty RAM.
2. BIOS or System Information (msinfo32).

Chapter Summary

- Systems use synchronous dynamic RAM (SDRAM) for primary memory. Static RAM (SRAM) is used for L1, L2, and L3 cache. Common versions of SDRAM are DDR, DDR2, and DDR3.
- Memory comes on circuit cards called DIMMs for desktop computers and SODIMMs for laptops. DIMMs and SODIMMs come in different sizes for different DDR versions.
- Dual-channel and triple-channel RAM provide additional 64-bit paths for transferring data to and from RAM. When installing multichannel DIMMs, install matched sets in the same bank. Banks are normally the same color. On an Intel dual-channel motherboard, Bank 0 includes slots 1 and 3.
- A triple-channel motherboard has six slots for RAM, and RAM should be purchased in matched sets of three DIMMs.
- The speed of RAM is tied directly to the clock. The formula to calculate DDR3 RAM speed is: Clk × 4 × 2 × 64 / 8. For a 200-MHz clock, the speed is 200 × 4 × 2 × 64 / 8, or 12,800 MB/s.
- The DDR3 standard name is derived from the clock × 8. For a 200-MHz clock, the DDR3 standard name is DDR3-1600. The module name is derived from the overall speed. The DDR3 module name with a 200-MHz clock is PC3-12800.
- If matched DIMMs are not used and a bank includes different speed DIMMs, the bank will default to the slowest speed.
- CPUs come in 32-bit and 64-bit versions, referring to how many bits they use to address memory. If you want to use more than 4 GB of RAM, you need a 64-bit CPU and a 64-bit operating system.
- Multiple-core CPUs include more than one fully functioning processor, and the operating system views each core as a separate CPU. Intel uses hyper-threading, which allows each core to process two threads at a time, and each core using hyper-threading is treated as a separate CPU by the operating system.
- CPUs use fast static RAM (SRAM) as cache to improve processing. They commonly include L1, L2, and sometimes L3 cache. L1 is fastest and closest to the CPU, and L3 is slowest and farthest away. L1 is smallest, and L3 is the largest. When the CPU needs data, it looks in L1, then L2, and then L3.
- The speed of the CPU is based on the speed of the clock and a multiplier. It is usually listed as the multiplied speed, such as 3.4 GHz. Intel uses Turbo Boost and AMD uses Turbo Core to modify these speeds during operation.
- Most CPUs support virtualization. Intel refers to its support as VT-x, and AMD calls its support AMD-V. These settings can be enabled in BIOS on most systems.
- An integrated GPU refers to a graphics processor embedded within a CPU. AMD calls some of its integrated GPU chips APUs.

- Common Intel CPUs are Core i3, Core i5, and Core i7 series. Most Intel CPUs use LGA sockets. Common Intel sockets are: LGA 775 (Socket T), LGA 1366 (Socket B), LGA 2011 (Socket R), LGA 1156 (Socket H or H1), and LGA 1155 (Socket H2).
- Common AMD CPUs are Sempron, Athlon, and Phenom. Most AMD CPUs use PGA sockets, and common sockets are: Socket 940, AM2, AM2+, AM3, AM3+, FM1, and Socket F.
- CPUs are commonly kept cool with heat sinks and fans. When replacing a CPU, use thermal paste between the CPU and the heat sink. Liquid cooling is an advanced cooling practice.
- Hardware problems that can cause unexpected shutdowns and intermittent fails include overheating due to failed fans or inadequate ventilation, faulty power supply, or faulty RAM.
- Use a software memory tester to test RAM.

Chapter Review

Use the following questions to test your knowledge of the information in this chapter. The answers to these questions, and the explanations of why each answer choice is correct or incorrect, are located in the "Answers" section at the end of this chapter.

1. You are replacing two DDR3 DIMMs in an Intel dual-channel motherboard. Into which slots should you put them?

   A. Two different-colored slots

   B. Two identical-colored slots

   C. Separate banks

   D. Slots 1 and 4

2. You are shopping for replacement DDR3 RAM. Your system has a 400-MHz clock. What should you buy?

   A. PC3-400

   B. DDR3-400

   C. PC3-25600

   D. PC3-12800

3. An Intel CPU has two cores, but the operating system shows it has four CPUs. What feature allows this to happen?
   - **A.** Hyper-threading
   - **B.** HyperTransport
   - **C.** Dual-channel RAM
   - **D.** L2 cache
4. Of the following choices, which is fastest?
   - **A.** L1 cache
   - **B.** L2 cache
   - **C.** L3 cache
   - **D.** Triple-channel DDR3
5. Which of the following replaces the Intel Socket H?
   - **A.** LGA 775
   - **B.** LGA 1366
   - **C.** LGA 1156
   - **D.** LGA 1155
6. You are asked to troubleshoot a computer that is randomly rebooting or failing. Of the following choices, what hardware can cause these symptoms? (Choose all that apply.)
   - **A.** RAM
   - **B.** Fan
   - **C.** Power supply
   - **D.** Virus

Answers

1. **Correct Answer:** B
   - **A. Incorrect:** Different-colored slots indicate different banks.
   - **B. Correct:** Dual-channel RAM should be installed in the same bank, which is the same color on most motherboards.
   - **C. Incorrect:** If you place the RAM in different banks, it will be used as single-channel RAM instead of dual-channel RAM.
   - **D. Incorrect:** Slots 1 and 4 are always in different banks.
2. **Correct Answer:** C
   - **A. Incorrect:** PC3-400 indicates a clock speed of 50 MHz.
   - **B. Incorrect:** If the DDR3 name is used, it is identified as the clock times 8. 400 × 8 = 3,200, or DDR3-3200.
   - **C. Correct:** The calculation for DDR3 is Clk × 4 × 2 × 64 / 8. 400 MHz × 4 × 2 × 64 / 8 = 25,600, so it is PC3-25600.
   - **D. Incorrect:** PC3-12800 indicates a clock speed of 200 MHz.
3. **Correct Answer:** A
   - **A. Correct:** Hyper-threading is supported on Intel CPUs and allows each core to appear as two CPUs.
   - **B. Incorrect:** HyperTransport is used on AMD processors in place of a front side bus.
   - **C. Incorrect:** Dual channel RAM provides two paths to RAM, but it does not affect the CPU cores.
   - **D. Incorrect:** L2 cache is fast RAM stored on the CPU for improved performance, but it does not affect the CPU cores.
4. **Correct Answer:** A
   - **A. Correct:** L1 cache is a fast cache, close to the CPU.
   - **B. Incorrect:** L2 cache is slower than L1 cache.
   - **C. Incorrect:** L3 cache is slower than L1 and L2 cache.
   - **D. Incorrect:** Any type of DDR RAM is slower than L1, L2, or L3 cache.
5. **Correct Answer:** D
   - **A. Incorrect:** LGA 775 is Socket T and was replaced by Socket B.
   - **B. Incorrect:** LGA 1366 is Socket B.
   - **C. Incorrect:** LGA 1156 is Socket H.
   - **D. Correct:** The LGA 1155 is also known as Socket H2 and replaces Socket H or H1.

6. **Correct Answers:** A, B, C
   - **A. Correct:** Faulty RAM can cause these symptoms.
   - **B. Correct:** Failing or dirty fans can result in overheating problems, causing these symptoms.
   - **C. Correct:** A power supply providing varying voltages or voltages out of specifications can cause these symptoms.
   - **D. Incorrect:** Viruses can cause these types of symptoms, but a virus is software, not hardware.