

Chapter 1

Getting Started

In This Chapter

- ▶ What computer programming is all about
 - ▶ Understanding the software that enables you write programs
 - ▶ Revving up to use an integrated development environment
-

Computer programming? What's that? Is it technical? Does it hurt? Is it politically correct? Does Bill Gates control it? Why would anyone want to do it? And what about me? Can I learn to do it?

What's It All About?

You've probably used a computer to do word processing. Type a letter, print it out, and then send the printout to someone you love. If you have easy access to a computer, then you've probably surfed the Web. Visit a page, click a link, and see another page. It's easy, right?

Well, it's easy only because someone told the computer exactly what to do. If you take a computer right from the factory and give no instructions to this computer, the computer can't do word processing, the computer can't surf the Web, it can't do anything. All a computer can do is follow the instructions that people give to it.

Now imagine that you're using Microsoft Word to write the great American novel, and you come to the end of a line. (You're not at the end of a sentence, just the end of a line.) As you type the next word, the computer's cursor jumps automatically to the next line of type. What's going on here?

Well, someone wrote a *computer program* — a set of instructions telling the computer what to do. Another name for a program (or part of a program) is *code*. Listing 1-1 shows you what some of Microsoft Word's code may look like.

Listing 1-1: A Few Lines in a Computer Program

```
if (columnNumber > 60) {
    wrapToNextLine();
}
else {
    continueSameLine();
}
```

If you translate Listing 1-1 into plain English, you get something like this:

```
If the column number is greater than 60,
then go to the next line.
Otherwise (if the column number isn't greater than 60),
then stay on the same line.
```

Somebody has to write code of the kind shown in Listing 1-1. This code, along with millions of other lines of code, makes up the program called Microsoft Word.

And what about Web surfing? You click a link that's supposed to take you directly to Yahoo.com. Behind the scenes, someone has written code of the following kind:

```
Go to <a href=http://www.yahoo.com>Yahoo</a>.
```

One way or another, someone has to write a program. That someone is called a *programmer*.

Telling a computer what to do

Everything you do with a computer involves gobs and gobs of code. Take a CD-ROM with a computer game on it. It's really a CD-ROM full of code. At some point, someone had to write the game program:

```
if (person.touches(goldenRing)) {
    person.getPoints(10);
}
```

Without a doubt, the people who write programs have valuable skills. These people have two important qualities:

- ✔ They know how to break big problems into smaller step-by-step procedures.
- ✔ They can express these steps in a very precise language.

A language for writing steps is called a *programming language*, and Java is just one of several thousand useful programming languages. The stuff in Listing 1-1 is written in the Java programming language.

Pick your poison

This book isn't about the differences among programming languages, but you should see code in some other languages so you understand the bigger picture. For example, there's another language, Visual Basic, whose code looks a bit different from code written in Java. An excerpt from a Visual Basic program may look like this:

```
If columnNumber > 60 Then
    Call wrapToNextLine
Else
    Call continueSameLine
End If
```

The Visual Basic code looks more like ordinary English than the Java code in Listing 1-1. But, if you think that Visual Basic is like English, then just look at some code written in COBOL:

```
IF COLUMN-NUMBER IS GREATER THAN 60 THEN
    PERFORM WRAP-TO-NEXT-LINE
ELSE
    PERFORM CONTINUE-SAME-LINE
END-IF.
```

At the other end of the spectrum, you find languages like ISETL. Here's a short ISETL program, along with the program's output:

```
{x | x in {0..100} | (exists y in {0..10} | y**2=x)};
{81, 64, 100, 16, 25, 36, 49, 4, 9, 0, 1};
```

Computer languages can be very different from one another but, in some ways, they're all the same. When you get used to writing IF COLUMN-NUMBER IS GREATER THAN 60, then you can also become comfortable writing if (columnNumber > 60). It's just a mental substitution of one set of symbols for another.

From Your Mind to the Computer's Processor

When you create a new computer program, you go through a multistep process. The process involves three important tools:

- ✓ **Compiler:** A compiler translates your code into computer-friendly (human-unfriendly) instructions.
- ✓ **Virtual machine:** A virtual machine steps through the computer-friendly instructions.
- ✓ **Application programming interface:** An application programming interface contains useful prewritten code.

The next three sections describe each of the three tools.

Translating your code

You may have heard that computers deal with zeros and ones. That's certainly true, but what does it mean? Well, for starters, computer circuits don't deal directly with letters of the alphabet. When you see the word *Start* on your computer screen, the computer stores the word internally as 01010011 01110100 01100001 01110010 01110100. That feeling you get of seeing a friendly looking five-letter word is your interpretation of the computer screen's pixels, and nothing more. Computers break everything down into very low-level, unfriendly sequences of zeros and ones, and then put things back together so that humans can deal with the results.

So what happens when you write a computer program? Well, the program has to get translated into zeros and ones. The official name for the translation process is *compilation*. Without compilation, the computer can't run your program.

I compiled the code in Listing 1-1. Then I did some harmless hacking to help me see the resulting zeros and ones. What I saw was the mishmash in Figure 1-1.

The compiled mumbo jumbo in Figure 1-1 goes by many different names:

- ✓ Most Java programmers call it *bytecode*.
- ✓ I often call it a *.class file*. That's because, in Java, the bytecode gets stored in files named *SomethingOrOther.class*.
- ✓ To emphasize the difference, Java programmers call Listing 1-1 the *source code*, and refer to the zeros and ones in Figure 1-1 as *object code*.

Figure 1-1:
My
computer
understands
these zeros
and ones,
but I don't.

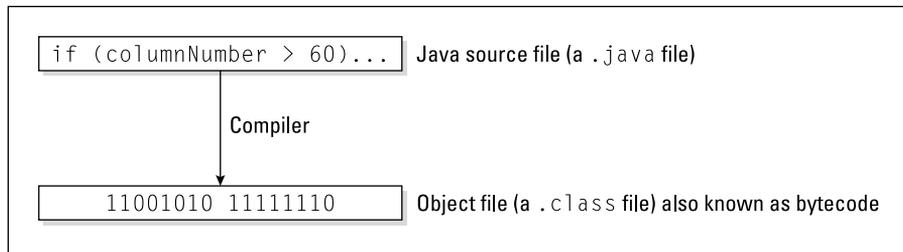
```

11001010 11111110 10111010 10111110 00000000 00000000
00000000 00101110 00000000 00010101 00001010 00000000
00000101 00000000 00010000 00010101 00000000 00000100
00000000 00010001 00001010 00000000 00000100 00000000
00010010 00000111 00000000 00010011 00000111 00000000
00010100 00000001 00000000 00000110 00111100 01101001
01101110 01101001 01110100 00111110 00000001 00000000
00000011 00101000 00101001 01010110 00000001 00000000
00000100 01000011 01101111 01100100 01100101 00000001
00000000 00001111 01001100 01101001 01101110 01100101
01001110 01110101 01101101 01100010 01100101 01110010
01010100 01100001 01100010 01101100 01100101 00000001
00000000 00001011 01100100 01101001 01110011 01110000
01101100 01100001 01111001 01010111 01101111 01110010
01100100 00000001 00000000 00000100 00101000 01001001
00101001 01010110 00000001 00000000 00001110 01110111
01110010 01100001 01110000 01010100 01101111 01001110
01100101 01111000 01110100 01001100 01101001 01101110
01100101 00000001 00000000 00010000 01100011 01101111
01101110 01110100 01101001 01101110 01110101 01100101
01010011 01100001 01101101 01100101 01001100 01101001
01101110 01100101 00000001 00000000 00001010 01010011
01101111 01110101 01110010 01100011 01100101 01100010

```

To visualize the relationship between source code and object code, see Figure 1-2. You can write source code, and then get the computer to create object code from your source code. To create object code, the computer uses a special software tool called a *compiler*.

Figure 1-2:
The
computer
compiles
source code
to create
object code.



Your computer's hard drive may have a file named `javac` or `javac.exe`. This file contains that special software tool — the compiler. (Hey, how about that? The word `javac` stands for “Java compiler!”) As a Java programmer, you often tell your computer to build some new object code. Your computer fulfills this wish by going behind the scenes and running the instructions in the `javac` file.

Running code

Several years ago, I spent a week in Copenhagen. I hung out with a friend who spoke both Danish and English fluently. As we chatted in the public park, I vaguely noticed some kids orbiting around us. I don't speak a word of Danish, so I assumed that the kids were talking about ordinary kid stuff.

Then my friend told me that the kids weren't speaking Danish. "What language are they speaking?" I asked.

"They're talking gibberish," she said. "It's just nonsense syllables. They don't understand English, so they're imitating you."

Now to return to present day matters. I look at the stuff in Figure 1-1, and I'm tempted to make fun of the way my computer talks. But then I'd be just like the kids in Copenhagen. What's meaningless to me can make perfect sense to my computer. When the zeros and ones in Figure 1-1 percolate through my computer's circuits, the computer "thinks" the thoughts in Figure 1-3.

Everyone knows that computers don't think, but a computer can carry out the instructions depicted in Figure 1-3. With many programming languages (languages like C++ and COBOL, for example), a computer does exactly what I'm describing. A computer gobbles up some object code, and does whatever the object code says to do.

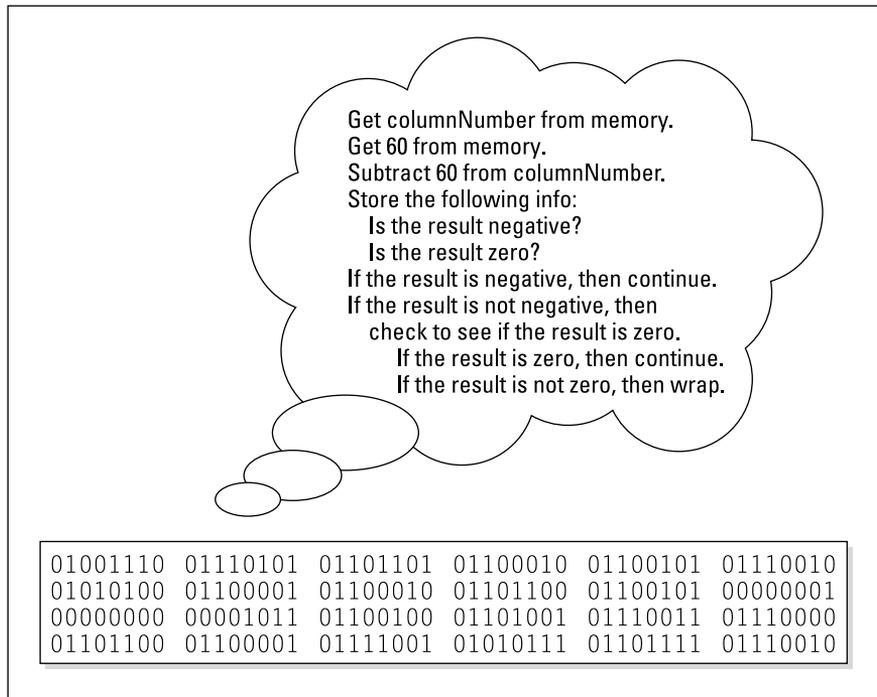


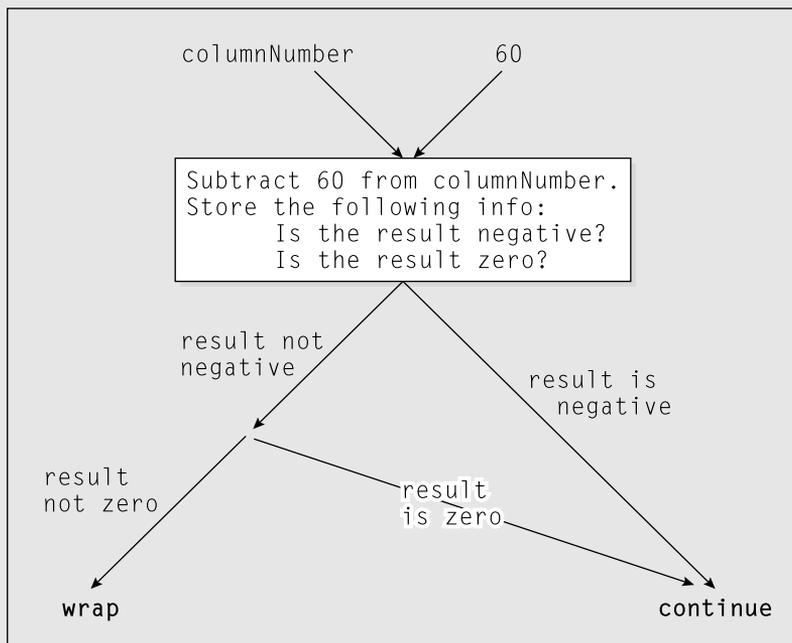
Figure 1-3:
 What the
 computer
 gleans from
 a bytecode
 file.

What is bytecode, anyway?

Look at Listing 1-1, and at the listing's translation into bytecode in Figure 1-1. You may be tempted to think that a bytecode file is just a cryptogram — substituting zeros and ones for the letters in words like `if` and `else`. But it doesn't work that way at all. In fact, the most important part of a bytecode file is the encoding of a program's logic.

The zeros and ones in Figure 1-1 describe the flow of data from one part of your computer to another. I've illustrated this flow in the following figure. But remember, this figure is just an illustration. Your computer doesn't look at this particular figure, or at anything like it. Instead, your computer reads a bunch of zeros and ones to decide what to do next.

Don't bother to absorb the details in my attempt at graphical representation in the figure. It's not worth your time. The thing you should glean from my mix of text, boxes, and arrows is that bytecode (the stuff in a `.class` file) contains a complete description of the operations that the computer is to perform. When you write a computer program, your source code describes an overall strategy — a big picture. The compiled bytecode turns the overall strategy into hundreds of tiny, step-by-step details. When the computer "runs your program," the computer examines this bytecode and carries out each of the little step-by-step details.



That's how it works in many programming languages, but that's not how it works in Java. With Java, the computer executes a different set of instructions. The computer executes instructions like the ones in Figure 1-4.

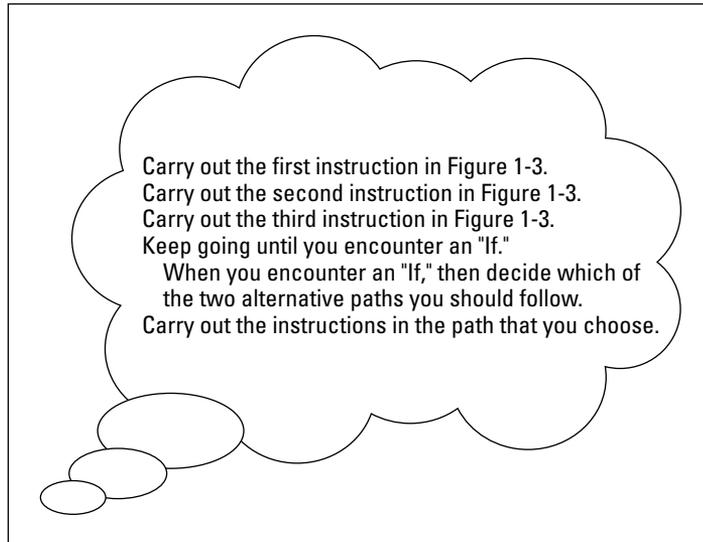


Figure 1-4:
How a
computer
runs a Java
program.

The instructions in Figure 1-4 tell the computer how to follow other instructions. Instead of starting with `Get columnNumber` from memory, the computer's first instruction is, "Do what it says to do in the bytecode file." (Of course, in the bytecode file, the first instruction happens to be `Get columnNumber` from memory.)

There's a special piece of software that carries out the instructions in Figure 1-4. That special piece of software is called the *Java virtual machine (JVM)*. The JVM walks your computer through the execution of some bytecode instructions. When you run a Java program, your computer is really running the Java virtual machine. That JVM examines your bytecode, zero by zero, one by one, and carries out the instructions described in the bytecode.

Many good metaphors can describe the Java virtual machine. Think of the JVM as a proxy, an errand boy, a go-between. One way or another, you have the situation shown in Figure 1-5. On the (a) side is the story you get with most programming languages — the computer runs some object code. On the (b) side is the story with Java — the computer runs the JVM, and the JVM follows the bytecode's instructions.

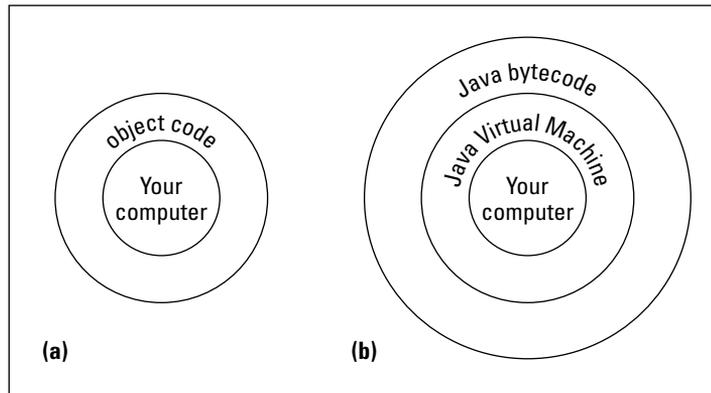


Figure 1-5:
Two ways
to run a
computer
program.



Your computer's hard drive may have a file named `java` or `java.exe`. This file contains the instructions illustrated previously in Figure 1-4 — the instructions in the Java virtual machine. As a Java programmer, you often tell your computer to run a Java program. Your computer fulfills this wish by going behind the scenes and running the instructions in the `java` file.

Code you can use

During the early 1980s, my cousin-in-law Chris worked for a computer software firm. The firm wrote code for word processing machines. (At the time, if you wanted to compose documents without a typewriter, you bought a “computer” that did nothing but word processing.) Chris complained about being asked to write the same old code over and over again. “First, I write a search-and-replace program. Then I write a spell checker. Then I write another search-and-replace program. Then, a different kind of spell checker. And then, a better search-and-replace.”

How did Chris manage to stay interested in his work? And how did Chris's employer manage to stay in business? Every few months, Chris had to reinvent the wheel. Toss out the old search-and-replace program, and write a new program from scratch. That's inefficient. What's worse, it's boring.

For years, computer professionals were seeking the Holy Grail — a way to write software so that it's easy to reuse. Don't write and rewrite your search-and-replace code. Just break the task into tiny pieces. One piece searches for a single character, another piece looks for blank spaces, a third piece substitutes one letter for another. When you have all the pieces, just assemble these pieces

to form a search-and-replace program. Later on, when you think of a new feature for your word processing software, you reassemble the pieces in a slightly different way. It's sensible, it's cost efficient, and it's much more fun.

The late 1980s saw several advances in software development, and by the early 1990s, many large programming projects were being written from prefab components. Java came along in 1995, so it was natural for the language's founders to create a library of reusable code. The library included about 250 programs, including code for dealing with disk files, code for creating windows, and code for passing information over the Internet. Since 1995, this library has grown to include more than 3,000 programs. This library is called the *API* — the *Application Programming Interface*.

Every Java program, even the simplest one, calls on code in the Java API. This Java API is both useful and formidable. It's useful because of all the things you can do with the API's programs. It's formidable because the API is so extensive. No one memorizes all the features made available by the Java API. Programmers remember the features that they use often, and look up the features that they need in a pinch. They look up these features in an online document called the *API Specification* (known affectionately to most Java programmers as the *API documentation*, or the *Javadocs*).

The API documentation describes the thousands of features in the Java API. As a Java programmer, you consult this API documentation on a daily basis. You can bookmark the documentation at the Sun Microsystems Web site and revisit the site whenever you need to look up something. But in the long run (and in the not-so-long run), you can save time by downloading your own copy of the API docs. (For details, see Chapter 2.)

Write Once, Run Anywhere™

When Java first hit the tech scene in 1995, the language became popular almost immediately. This happened in part because of the Java virtual machine. The JVM is like a foreign language interpreter, turning Java bytecode into whatever native language a particular computer understands. So if you hand my Windows computer a Java bytecode file, then the computer's JVM interprets the file for the Windows environment. If you hand the same Java bytecode file to my colleague's Macintosh, then the Macintosh JVM interprets that same bytecode for the Mac environment.

Look again at Figure 1-5. Without a virtual machine, you need a different kind of object code

for each operating system. But with the JVM, just one piece of bytecode works on Windows machines, Unix boxes, Macs, or whatever. This is called *portability*, and in the computer programming world, portability is a very precious commodity. Think about all the people using computers to browse the Internet. These people don't all run Microsoft Windows, but each person's computer can have its own bytecode interpreter — its own Java virtual machine.

The marketing folks at Sun Microsystems call it the *Write Once, Run Anywhere™* model of computing. I call it a great way to create software.

Your Java Programming Toolset

To write Java programs, you need the tools described previously in this chapter:

- ✔ **You need a Java compiler.** (See the section entitled, “Translating your code.”)
- ✔ **You need a Java virtual machine.** (See the section entitled, “Running code.”)
- ✔ **You need the Java API.** (See the section entitled, “Code you can use.”)
- ✔ **You need the Java API documentation.** (Again, see the “Code you can use” section.)

You also need some less exotic tools:

- ✔ **You need an editor to compose your Java programs.**

Listing 1-1 contains part of a computer program. When you come right down to it, a computer program is a big bunch of text. So to write a computer program, you need an *editor* — a tool for creating text documents.

An editor is a lot like Microsoft Word, or like any other word processing program. The big difference is that an editor adds no formatting to your text — no bold, no italic, no distinctions among fonts. Computer programs have no formatting whatsoever. They have nothing except plain old letters, numbers, and other familiar keyboard characters.

- ✔ **You need a way to issue commands.**

You need a way to say things like “compile this program” and “run the Java virtual machine.”

Every computer provides ways of issuing commands. (You can double-click icons or type verbose commands in a Run dialog box.) But when you use your computer’s facilities, you jump from one window to another. You open one window to read Java documentation, another window to edit a Java program, and a third window to start up the Java compiler. The process can be very tedious.

In the best of all possible worlds, you do all your program editing, documentation reading, and command issuing through one nice interface. This interface is called an *integrated development environment* (IDE).

A typical IDE divides your screen’s work area into several panes — one pane for editing programs, another pane for listing the names of programs, a third pane for issuing commands, and other panes to help you compose and test programs. You can arrange the panes for quick access. Better yet, if you change the information in one pane, the IDE automatically updates the information in all the other panes.

Some fancy environments give you point-and-click, drag-and-drop, plug-and-play, hop-skip-and-jump access to your Java programs. If you want your program to display a text box, then you click a text box icon and drag it to the workspace on your screen.

Figure 1-6 illustrates the use of a drag-and-drop IDE. In Figure 1-6, I create a program that displays two images, two text fields, and two buttons. To help me create the program, I use the Eclipse IDE with the Jigloo graphical plug-in. (For a taste of Eclipse, visit www.eclipse.org. For more info on the neat Jigloo graphical user interface builder, check out www.cloudgarden.com.)

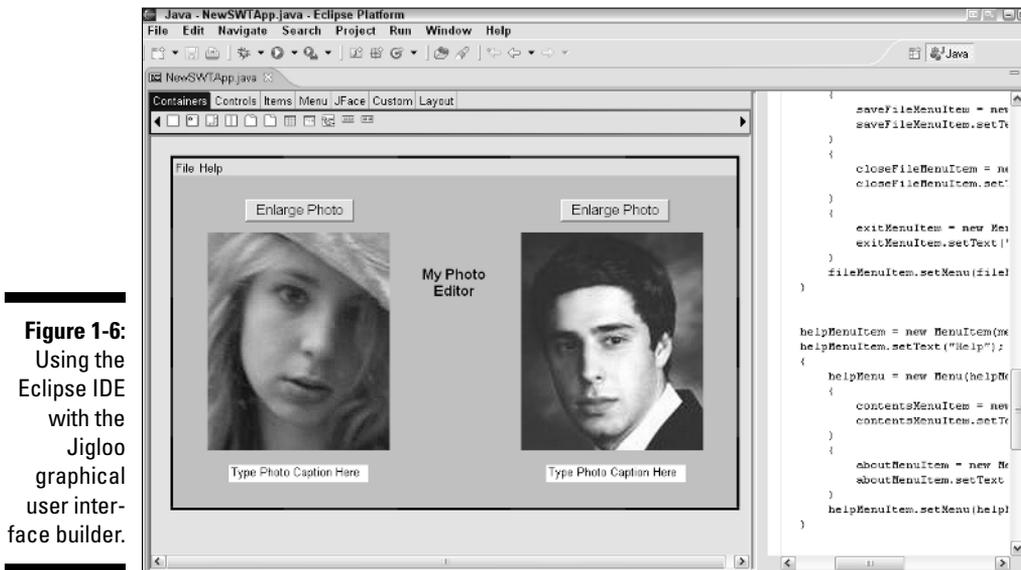


Figure 1-6: Using the Eclipse IDE with the Jigloo graphical user interface builder.

An IDE helps you move seamlessly from one part of the programming endeavor to another. With an IDE, you don't have to worry about the mechanics of editing, compiling, and running a Java virtual machine. Instead, you can worry about the logic of writing programs. (Wouldn't you know it? One way or another, you always have something to worry about!)

What's already on your hard drive?

You may already have some of the tools you need for creating Java programs. Here are some examples:

- ✔ **Most versions of Windows come with a Java virtual machine.** Look for a file named `java.exe` in your `\windows\system32` directory.
- ✔ **Most computers running Mac OS X come with a Java compiler, a Java virtual machine, and a Java API.**
- ✔ **Some IDEs come with their own Java tools.** For example, when you buy the Borland JBuilder IDE you get a compiler, a Java virtual machine, and a copy of the Java API. When you download the free Eclipse IDE you get a Java compiler, but no Java virtual machine and no Java API.

You may already have some Java tools, but your tools may be obsolete. This book's examples use a relatively new version of Java — a version released in September 2004. Even computers and software sold in 2005 may not be up to date with the latest Java features. So if you use the tools that come with your computer, or if you use a commercial product's software tools, some of this book's examples may not run.

The safest bet is to download tools afresh from the Sun Microsystems Web site. To get detailed instructions on doing the download, see Chapter 2.



Many of this book's examples don't run on "older" versions of Java, and by "older" I mean versions created before the fall of 2004. If you have trouble running the programs in this book, check to make sure that your version of Java is numbered 5.0, 5.1, or something like that. Older versions (with version numbers like 1.4 or 1.4.2) just don't cut the muster.

JCreator

The programs in this book work with any IDE that can run Java 5.0. You can even run the programs without an IDE. But to illustrate the examples in this book, I use JCreator LE (Lite Edition). I chose JCreator LE over other IDEs for several reasons:

- ✔ JCreator LE is free.
- ✔ Among all the Java IDEs, JCreator represents a nice compromise between power and simplicity.
- ✔ Unlike some other Java IDEs, JCreator works with almost any version of Java, from the ancient version 1.0.2 to the new-and-revolutionary version 5.0.
- ✔ JCreator LE is free. (It's worth mentioning twice.)

This book's Web site has a special edition of JCreator LE — a version that's customized especially for *Beginning Programming with Java For Dummies*, 2nd Edition readers! For details on downloading and installing the special edition of JCreator, see Chapter 2.

JCreator runs only on Microsoft Windows. If you're a Unix, Linux, or Macintosh user, please don't be offended. All the material in this book applies to you, too. You just have to use a different IDE. My personal recommendations include Eclipse and Netbeans. For details, visit this book's Web site at <http://www.dummies.com/go/bpjavafd>.